440-TP-006-002

Production Topologies: A Trade-off Study Analysis for the ECS Project

Technical Paper

August 1995

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEER

N. Prasad /s/	8/25/95
Narayan S. Prasad, Scientist/Engineer EOSDIS Core System Project	Date

SUBMITTED BY

Parag N. Ambardekar /s/	8/25/95
Parag Ambardekar, Release A Manager EOSDIS Core System Project	Date

Hughes Information Technology Corporation
Upper Marlboro, Maryland

This page intentionally left blank.

Abstract

Within the current processing architecture framework, chains of processing resources are pooled or dedicated to support Ad Hoc Working Group on Production (AHWGP) requirements. These equipment chains previously called "string" within the SDPS SDS (also called cluster or subnetwork) are composed of hardware laid out to perform processing, reprocessing, science software integration and test, and backup processing, while minimizing impacts on communications and other staging infrastructure wherever possible. Each individual cluster may be designed to support the unique processing and reprocessing requirements of product generation tasks assigned to it. The concept of processing clusters has yielded a series of recommended physical (not logical) processing topologies that can impact hardware requirements, overall performance, network capacity, staging storage, product generation throughput, etc.

In the first part of this trade-off study analysis issued at PDR (Ref number: 440-TP-006-001), four cluster optimization alternatives were identified to do production processing. They are: a) One instrument's products per cluster, b) One instrument's products per cluster except for selected products requiring major processing resources, c) Multiple instruments' products on any cluster, and d) Any instruments' products on any cluster that can support it.In the second phase of this trade study, a static and a more detailed dynamic analysis of AHWGP data is performed for the third quarter of 1999 (Release B/C) time period for the first optimization alternative, namely "one instrument's products per cluster". Four DAACs (LaRC, GSFC, EDC and MSFC) have been chosen for this study. The January 1995 AHWGP baseline is used as input for both static and dynamic analyses. The ECS System Performance Model is used to dynamically simulate processing. Due to limitations with the current version of the ECS Systems Performance Model in simulating all cluster optimization alternatives, only one production topology is analyzed dynamically. Therefore, recommendations to the most appropriate topology at each site is deferred until further analysis with more topologies is completed. This paper is intended to provide a framework and identify factors that must be evaluated for each of the cluster optimization alternatives presented in the PDR release of the document.

Keywords: Topologies, cluster optimization alternatives, ECS Systems Performance Model, string, subnetwork, processing, dynamic analysis, static analysis, AHWGP

This page intentionally left blank.

Contents

Abstract

1. Introduction

1.1 Trade description 1

1.2 Scope 1

1.3 Organization 2

1.4 Acknowledgments 2

1.5 Review and Approval 2

1.6 Applicable and Reference Documents 3

2. Executive Summary

2.1 Major Analysis/implementation Alternatives 4

2.2 Analysis Summary 5

3. Background

3.1 Background of Data Processing Subsystem 7

3.2 Background of ECS Systems Performance Model 7

3.3 AHWGP Data 9

4. Cluster Optimization Alternatives

4.1 Introduction 10

4.2 One instrument's products per cluster 11

4.2.1 Physical view 11

4.3 One instrument's products per cluster except for selected products requiring major processing resources 12

4.3.1 Physical view 12

4.4 Multiple instruments' products on any cluster 13
4.4.1 Physical view 13
4.5 Any instruments products on any cluster that can support it 13
4.5.1 Physical view 13

5. Processing Requirements

5.1 Introduction 15
5.2 Release A 15
5.2.1 Processing Requirements for LaRC at Release A 15
5.2.2 Processing Requirements for MSFC at Release A 18
5.3 Release B 19
5.3.1 Processing Requirements for LaRC at Release B 19
5.3.2 Processing requirements by instrument at GSFC 25
5.3.3 Processing requirements by instrument at EDC 30
5.3.4 Processing Requirements at MSFC 33

6. Analysis By Cluster Optimization Alternative 34

6.1 General Assumptions 34
6.2 Cluster Optimization Alternatives Drivers 34
6.3 Analysis by optimization alternative 35
6.3.1 One instrument's products per cluster 35
6.3.2 One instrument's products per cluster except for selected products requiring major processing resources 85
6.3.3 Multiple instruments' products per cluster 86
6.3.4 Any instrument's products on any cluster that can support it 86

7. Summary and Recommendations

Figures

3.2-1. Top-level Module of ECS Systems Performance Model 8
6.3-1a Schematic of Data Movement at CPU for Host Attached Storage 36
6.3-2a. CERES Processing Resource Usage - Number of CPUs (300 MFLOPS) 58
6.3-2b. CERES Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 59
6.3-3a. CERES Processing Resource Usage - Number of CPUs (900 MFLOPS) 60
6.3-3b. CERES Processing Resource Usage - Processing Disk Capacity (900 MFLOPS) 61
6.3-4a. MISR Processing Resource Usage - Number of CPUs (300 MFLOPS) 63
6.3-4b. MISR Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 64
6.3-5a. MISR Processing Resource Usage - Number of CPUs (900 MFLOPS) 65
6.3-5b. MISR Processing Resource Usage - Processing Disk Capacity (900 MFLOPS) 66
6.3-6a. MOPITT Processing Resource Usage - Number of CPUs (300 MFLOPS) 68
6.3-6b. MOPITT Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 69
6.3-7a. MODIS (GSFC) Processing Resource Usage - Number of CPUs (300 MFLOPS) 72
6.3-7b. MODIS (GSFC) Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 73
6.3-8a. MODIS (GSFC) Processing Resource Usage - Number of CPUs (900 MFLOPS) 74
6.3-8b. MODIS (GSFC) Processing Resource Usage - Processing Disk Capacity (900 MFLOPS) 75
6.3-9a. ASTER Processing Resource Usage - Number of CPUs (300 MFLOPS) 78
6.3-9b. ASTER Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 79
6.3-10a. ASTER Processing Resource Usage - Number of CPUs (900 MFLOPS) 80
6.3-10b. ASTER Processing Resource Usage - Processing Disk Capacity (900 MFLOPS) 81
6.3-11a. LIS Processing Resource Usage - Number of CPUs (300 MFLOPS) 83
6.3-11b. LIS Processing Resource Usage - Processing Disk Capacity (300 MFLOPS) 84

Tables

5.2-1. CERES Processing Requirements Summary for Release A 16

5.2-2. Summary of Definitions...18

5.2-3. LIS Processing Requirements Summary for Release A...19

5.3-1. ACRIM Processing Requirements Summary 19

5.3-2. CERES Processing Requirements Summary for Release B 20

5.3-3. MISR Processing Requirements Summary...23

5.3-4. MOPITT Processing Requirements Summary...24

5.3-5. COLOR Processing Requirements Summary...25

5.3-6. MODIS (GSFC) Processing Requirements Summary...27

5.3-7. ASTER Processing Requirements Summary...31

5.3-8. MODIS (EDC) Processing Requirements Summary...32

5.3-9. LIS Processing Requirements Summary...33

6.3-1a. CERES Daily Average Processor Requirements as a Function of Duty Cycle for Release A 38

6.3-1b. CERES Daily Average Processor Requirements as a Function of Duty Cycle for Release B...38

6.3-2a. CERES Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release A LaRC...............................40

6.3-2b. CERES Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC 41

6.3-3. MISR Daily Average Processor Requirements as a Function of Duty Cycle for Release B...43

6.3-4. MISR daily average theoretical I/O bandwidth at CPU for Release B at LaRC 44

6.3-5. MOPITT Daily Average Processor Requirements as a Function of Duty Cycle 44

6.3-6. MOPITT Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC 45

6.3-7. ACRIM Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC 45

6.3-8. LIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Releases A and B at MSFC..............46

6.3-9. ASTER Daily Average Processor Requirements as a Function of Duty Cycle for Release B....46

6.3-10. ASTER Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at EDC....47

6.3-11. MODIS (EDC) Daily Average Processor Requirements as a Function of Duty Cycle for Release B....48

6.3-12. MODIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at EDC....49

6.3-13. COLOR Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC....49

6.3-14. COLOR Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC....50

6.3-15. MODIS Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC....51

6.3-16. MODIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at GSFC....52

6.3-17b CERES Process Completion Times for 900-MFLOPS....56

6.3-18a. CERES CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)....61

6.3-18b. CERES CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)....62

6.3-19a. MISR Process Completion Times (300 MFLOPS)....62

6.3-19b. MISR Process Completion Times (900 MFLOPS)....62

6.3-20a. MISR CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)....66

6.3-20b. MISR CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)....66

6.3-21. MOPITT Process Completion Times (300 MFLOPS)....67

6.3-22. MOPITT CPU and Staging disk capacity from ECS Systems Performance Model (300 MFLOPS)....69

6.3-23a. MODIS (GSFC) Process Completion Times (300 MFLOPS)....70

6.3-23b. MODIS (GSFC) Process Completion Times (900 MFLOPS)....71

6.3-24a. MODIS (GSFC) CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)....75

6.3-24b. MODIS (GSFC) CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)....75

6.3-25a. ASTER Process Completion Times (300 MFLOPS)....76

6.3-25b. ASTER process completion times (900 MFLOPS)....77

6.3-26a. ASTER CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)....81

6.3-26b. ASTER CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)....81

6.3-27. LIS Process Completion Times (300 MFLOPS)....82

6.3-29. Processing <--> Data Handler Network Flow by DAAC....85

Abbreviations and Acronyms

1. Introduction

1.1 Trade description

Within the current processing architecture framework, chains of processing resources are pooled or dedicated to support Ad Hoc Working Group on Production (AHWGP) requirements. These equipment chains previously called "string" within the SDPS SDS (also called cluster or subnetwork) are composed of hardware laid out to perform processing, reprocessing, science software integration and test, and backup processing, while minimizing impacts on communications and other staging infrastructure wherever possible. Each individual cluster may be designed to support the unique processing and reprocessing requirements of product generation tasks assigned to it. The concept of processing clusters has yielded a series of recommended physical (not logical) processing topologies that can impact hardware requirements, overall performance, network capacity, staging storage, product generation throughput, etc. This trade examines the pros and cons of distributing processing tasks from one or more instruments across one or more processing clusters. Recommendations will be made for the most cost effective way of distributing processing to maximize throughput, minimize data movement, and provide and retain the flexibility to evolve with changing processing requirements.

1.2 Scope

In the first part of this trade-off study analysis issued at PDR (Ref number: 440-TP-006-001), four cluster optimization alternatives were identified to do production processing. In the second phase of this trade study, a static and a more detailed dynamic analysis of AHWGP data is performed for the third quarter of 1999 (Release B/C) time period for the first optimization alternative, namely "one instrument's products per cluster". Four DAACs (LaRC, GSFC, EDC and MSFC) are chosen for this study. The January 1995 AHWGP baseline is used as input for both static and dynamic analyses. The ECS System Performance Model is used to dynamically simulate processing. The LaRC DAAC will process CERES, MISR and MOPITT instrument data. Because CERES and MISR have large processing requirements and I/O loads, and together with MOPITT have external data dependencies (e.g. MODIS products), they will provide insight into the architectural constraints, if any for the production topologies considered. Similarly MODIS at GSFC will provide insight into the kind of topologies required at that site. The ASTER and MODIS (L3 processing was not included) processing at EDC is also evaluated based on the data provided by the AHWGP. Because of the small processing requirements for the LIS instrument, it will not be necessary to consider various production topologies at MSFC. Nonetheless, both a static and dynamic analyses is performed to understand capacity requirements.

Due to limitations with the current version of the ECS Systems Performance Model in simulating all cluster optimization alternatives (see Section 4.2), only one production topology is analyzed dynamically. Therefore, recommendations to the most appropriate topology at each site is deferred until further analysis with more topologies is performed. This paper is intended to provide a framework and identify factors that must be evaluated for each of the cluster optimization alternatives presented in the PDR release of the document.

1.3 Organization

This paper is organized as follows:

An executive summary provides an outline of major analysis, implementation alternatives, and preliminary results reported in this technical paper. Section 3 gives a background of the Data Processing Subsystem, the ECS System Performance Model and AHWGP data. The various cluster optimization alternatives are discussed in Section 4. Section 5 provides the processing requirements for each instrument for both Releases and A and B. Section 6 details the analysis performed based on the AHWGP data for four Release B DAACs (EDC, GSFC, LaRC and MSFC). Detailed results from the dynamic analysis using the System Performance Model is presented for one of the cluster optimization alternatives under consideration. Advantages and disadvantages of the various optimization alternatives are also listed. Section 7 draws conclusions based on the analysis. Acronyms used throughout the text are listed in the Abbreviations and Acronyms List.

1.4 Acknowledgments

The assistance given by the ECS System Performance Modeling team (Bob Howard, Mike Theobald and Rajesh Dharia) is sincerely acknowledged.

1.5 Review and Approval

Questions regarding technical information contained within this paper should be addressed to the following ECS and/or GSFC contacts:

- ECS Contact
  Narayan Prasad, Scientist/Engineer
  (301)-925-0467
  nprasad@eos.hitc.com

- GSFC Contact
  Steve Kempler, PDPS Manager
  (301)-286-7766
  steven.j.kempler@gsfc.nasa.gov

Questions concerning distribution or control of this document should be addressed to:

Data Management Office
The ECS Project Office
Hughes Applied Information Systems
1616 McCormick Dr.
Upper Marlboro, MD 20774

1.6 Applicable and Reference Documents

1. Science Data Processing Segment (SDPS) Segment Design Specifications for the ECS Project, 305-CD-002-001.
2. Systems Performance Model for the ECS Project, 241-TP-001-001.
3. Investigator guide to estimating EOS data production. Bruce Barkstrom and members of the Ad Hoc Working Group on Production (unpublished).
4. Network Attached Storage Concepts and Industry Survey, 440-TP-009-001.
5. Trade-off Studies Analysis Data for the ECS Project, 211-CD-001-002.

2. Executive Summary

2.1 Major Analysis/implementation Alternatives

In the second phase of this trade (further analysis to be performed by Release B group), a static and a detailed dynamic analyses of the AHWGP data is made for the Release B/C time period during the third quarter of 1999 (epoch "k") for four Release B DAACs (EDC, GSFC, MSFC and LaRC). This time period is chosen because it is when maximum capacity requirements are exercised. The AHWGP data are input into the ECS Systems Performance Model. Only Standard Processing is considered for this analysis. Combinations of processing scenarios with various processing topologies will be explored to serve as a guideline to determine optimized configurations of processing hardware. The calculations are based on the January 1995 AHWGP baseline. The AHWGP numbers may be refined if processing requirements change.

It is important to describe the concept of clusters from an operational viewpoint. Clusters are physical optimizations that do not prevent wholesale pooling for processing or reprocessing campaigns. The concept of processing clusters, current performance requirements, the resource pooling and dedication trade analysis at SDR has yielded different candidate cluster formations or cluster formation optimization alternatives which optimize different selection criteria (e.g. communications, staging, RMA (Reliability, Maintainability, Availability), ease of operations, management and control). It should be emphasized that these alternatives are fully configurable on a case-by-case basis, making it flexible to handle changing requirements or by release. The overall Planning and Data Processing architecture is built on the concepts of resource pooling regardless of the physical network layout. The different subnetwork/string/cluster formation alternatives are:

- *One instrument's products per cluster* -- In this option, each instrument has a processing cluster, consisting of one or more compute servers dedicated to the production of its products. A detailed static and dynamic analyses is performed for this optimization alternative.
- *One instrument's products per cluster except for selected products requiring major processing resources* -- This option is identical to the first option except that at certain times when processing resources for certain products of that instrument exceeds the maximum allowable resources of that processing cluster, then processing of that product can shift to the cluster that has the resources to support it. Dynamic analysis is deferred until the ECS Systems Performance Model attains sufficient maturity to handle this alternative.
- *Multiple instruments' products on any cluster* -- This may apply to conditions whereby instruments with interdependent processing may be collocated. This situation does not apply to the LaRC scenario. There is no interdependency among products from the three instruments. Also, the MODIS products that each instrument requires are different.

Dynamic analysis for the other DAACs is deferred until the ECS Systems Performance Model attains sufficient maturity to handle this alternative.

- *Any instruments' products on any cluster that can support it* **--** This option is a mix-and-match situation. The processing load will determine the cluster where a particular instrument's products will be processed. This option needs a full scale dynamic analysis. It will be explored after the ECS Systems Performance Model attains sufficient maturity to handle this alternative.

The allowable optimized cluster formation alternatives may:

- be more than one at any given site. This may be the case when a site handles many instruments, and the processing requirements among instruments show large spread (e.g. MOPITT's processing requirements are small compared to CERES and MISR);
- differ from one site to another because each DAAC handles different instruments;
- differ with each release when significantly more complexities are introduced into the system and processing requirements increase.

2.2 Analysis Summary

Physical cluster optimization, on a site-by-site basis for Release A, is not a major concern due to the small numbers and scale of the physical equipment currently envisioned for activation at that time. Release A implementations predicted for operations for LaRC and MSFC involve mid-performance (predicted) LANs and only two physical science processors within the SPRHW CI [1]. The GSFC configuration, which does not support processing operations, involves one or a small number of compute resources at a maximum. Thus, single physical subnetworks can be used (with the proper backup for RMA concerns) to couple the processing resources with primary ingest and Data Server resources, for example. The driver on selecting more than one subnetwork will be the actual throughput rates required, as opposed to operational or mission requirements which form the real basis of the implementation alternatives summarized earlier. It is expected that physical subnetwork optimization will be a larger issue for Releases B and beyond. Therefore this analysis provides a framework and approach for a more detailed analysis required for Release B hardware implementation.

The key recommendation is that multiple strings/cluster/subnetwork formation alternatives and selection criteria be allowed both between DAAC sites and within. One implementation alternative for all sites and all releases is not recommended. This will permit subnetworks of ECS resources to be tuned to meet the primary needs of the DAAC site, but will still allow view of the resources (through planning and production management) as a single processing pool or series of subpools.

The first option (one instrument's products per cluster) is a more natural way of doing data processing. Based on the requirements from the AHWGP, the three instruments can be each assigned to independent clusters. An obvious disadvantage of this set up is that the processing resources on a cluster may not be fully utilized, while a backlog can occur on another. As an example, a static analysis of the AHWGP data has yielded the following. MOPITT L1 and L2 processes are activated only once a day, while L3 processes are activated only weekly. With daily average MFLOP requirements for MOPITT at less than 20 (based on raw numbers), it is a good candidate for sharing resources with MISR or CERES.

A static analysis of the AHWGP data has yielded a daily average MFLOP requirements for MISR for this time period to be 3455, with an "I/O bandwidth at CPU" of 19.3 MB/s. Since MISR requires large volumes of data to be staged (4.2 MB/s) and destaged (2.2 MB/s) (due to the large number of activations per day), it appears that distributing (according to optimization alternatives discussed earlier) MISR processing can increase network traffic, which in turn can degrade overall performance.

The CERES Data Processing Subsystems 4 (determine cloud properties, and top of the atmosphere and surface fluxes) and 5 (compute surface and atmospheric fluxes) take up more than 90% of the total CERES MFLOP requirements (986 and 1782 MFLOPS for Subsystems 4 and 5, respectively). The daily average I/O bandwidth at CPU for Subsystems 4 and 5 are 3.13 MB/s and 0.7 MB/s, respectively. With relatively low I/O bandwidth at CPU, other CERES Subsystems (excluding 4 and 5) could share resources with MOPITT.

Similarly the ECS Systems Performance Model is used under the first configuration (one instrument's products per cluster) for other Release B DAACs. The volume of the staging disk and the optimal number of processors are analyzed for each instrument.

3. Background

3.1 Background of Data Processing Subsystem

The Data Processing Subsystem (DPS) consists of three hardware CIs namely: 1) Science processing (SPRHW), 2) Algorithm Integration and Test (AITHW, and 3) Quality Assessment and Monitoring (AQAHW). It is responsible for managing, queuing and executing processes on a specified set of processing resources at each DAAC site. The science processing resources can be a chain of processing resources known as "clusters". They are self-contained processing resources based on a set series of alternatives for selection. They may also imply chains of processing, I/O and staging resources configured to deal with unique processing requirements to which they are allocated. This does not imply that the only use of that processing cluster, or a specific compute server on that cluster, is for only one specific class of instrument algorithms alone. It should be emphasized that cluster formations are configurable on a case-by-case basis, making it fully flexible to handle changing requirements by release. Clusters can be used for both processing and reprocessing. One or more data servers stage data on to the relevant working storage pool allocated to the processing clusters. Separate cluster resources are allocated for algorithm integration and test. This cluster called the Test and Backup cluster is configured with a like complement of processing, I/O and attached staging resources. Each processing cluster is supported by the Test and Backup cluster. The cluster topology provides a "fail soft" environment almost by its very nature.

3.2 Background of ECS Systems Performance Model

The ECS Systems Performance Model is a Block Oriented Network Simulation (BONeS) model [2] which is used in conjunction with the AHWGP data to simulate processing. Figure 3.2-1 contains the top-level module of BONeS with representation of ECS Subsystems. A brief description of the model components follow:

BONeS is a discrete-event simulation tool for analysis and design of communication networks and distributed processing systems. The components of a distributed processing system (including the networks) are represented by nodes. Nodes have resources associated with them which get allocated as events request them. Standard production of instrument data within DPS is simulated by the Processing module, in conjunction with Event Driven Scheduler and the Data Handler. The Data Handler is the model's representation of the Data Server design. It is responsible for storing and retrieving data from the permanent archive, for routing data to the requesting subsystems, and for managing tiered storage resources. The scheduler monitors the availability of data, requests data to be staged from the data Handler to Processing, routes newly created data to the appropriate data handler or processing pool, and initiates execution of a process when all required inputs are present. The Ingest module emulates behavior of the Ingest subsystem: acceptance of data from external systems and users, rolling storage of L0 instrument data, etc. are handled here. The Distribution module simulates network and media distribution of data to users.

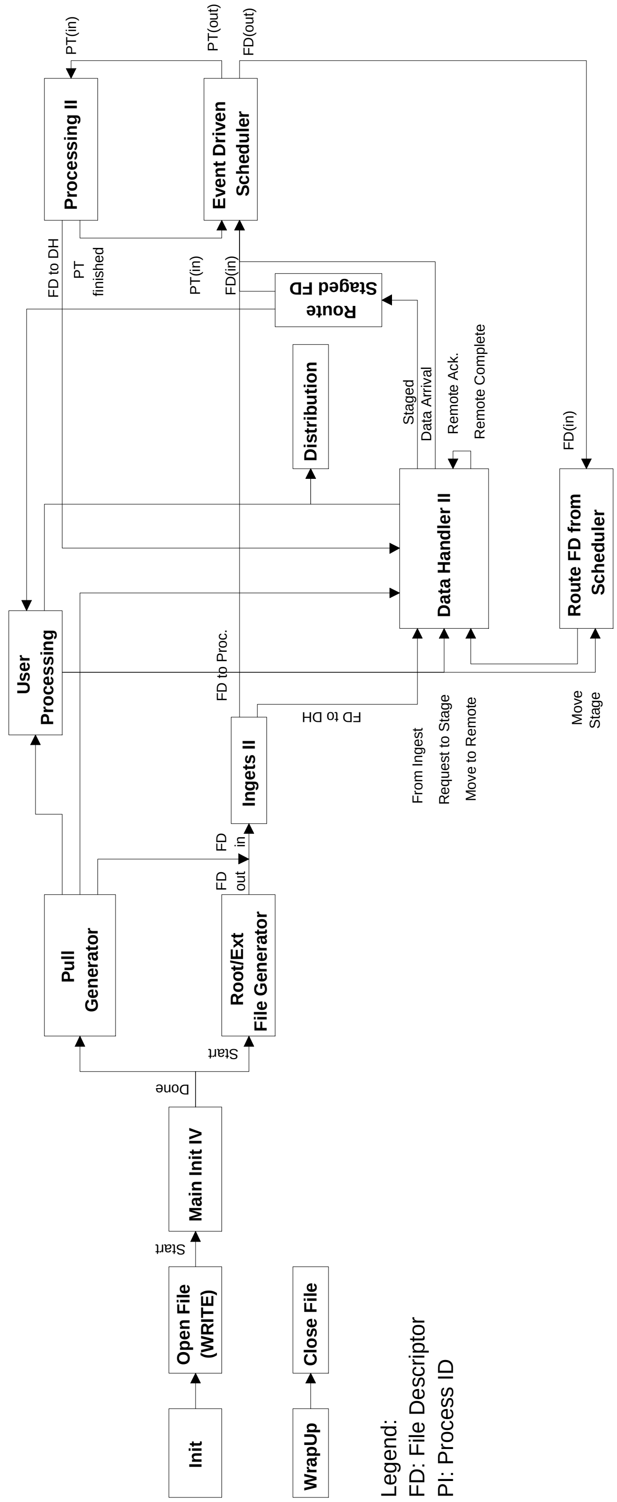

Figure 3.2-1 Top-level Module of ECS Systems Performance Model

During a simulation, the program collects data at selected points in the model network using a variety of probes. The model simulation is "resource constrained" in that the nodes in the model are specified to correspond to a particular system configuration. For example, the number of processors in a processing cluster may be constrained. The model will then determine for the constrained number of processors the time required to handle the data volume for normal operations. The performance of this configuration is measured by the simulation. The model is currently tested to verify operations within the entire suite of AHWGP data. The model will incrementally be made sophisticated to support an array of planned experiments and design tradeoffs.

3.3 AHWGP Data

The Ad Hoc Working Group on Production [3] is represented by members of ASTER, CERES, LIS, MISR, MODIS, MOPITT and other instruments. The AHWGP was formed to produce a reliable estimate of the computer and network resources required to support ECS data production, and to provide ECS modelers and the Project information on data production plans. More specifically data from the AHWGP includes data products and their sizes, names of processes, number of activations and their activation scenarios, an estimate of the CPU processing capacity for each process, staging disk storage, number of file transfers and their sizes. The January 1995 version of the AHWGP data was used for this study although a June 1995 version has been baselined recently. The June 1995 baseline currently has not been incorporated into the ECS Systems Performance Model.

4. Cluster Optimization Alternatives

4.1 Introduction

Clusters are physical optimizations that do not prevent wholesale pooling for processing or reprocessing campaigns. The concept of processing clusters, current performance requirements, the resource pooling and dedication trade analysis at SDR (unpublished) have yielded different candidate cluster formations or cluster formation optimization alternatives which optimize different selection criteria (e.g. communications, staging, RMA, ease of operations, management and control). It should be emphasized that these alternatives are configurable on a case-by-case basis, making it fully flexible to handle changing requirements or by release. The overall Planning and Data Processing architecture is built on the concepts of resource pooling regardless of the physical network layout. Figure 4.1-1 illustrates a generic star topology of a cluster formation. This topology is presented only as an example and does not imply that DAAC hardware will be configured this way. The DAAC topologies are driven by DAAC unique requirements (see Section A.1-A-3 in the DAAC Unique Appendices for operational sites [1]).

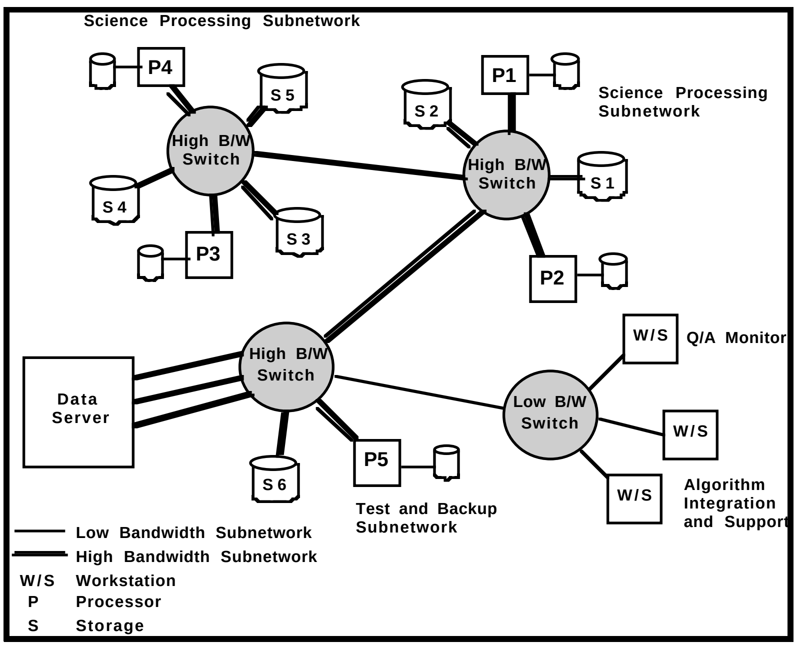

Figure 4.1-1 Generic Star Topology of a Cluster

The allowable optimized cluster formation alternatives may:

- be more than one at any given site. This may be the case when a site handles many instruments, and the processing requirements among instruments show large spread (e.g. MOPITT's processing requirements are small compared to CERES and MISR);
- differ from one site to another because each DAAC handles different instruments;
- differ with each release when significantly more complexities are introduced into the system and processing requirements increase.

The different subnetwork/string/cluster formation alternatives are discussed in the following sections:

4.2 One instrument's products per cluster

In this option, each instrument has a processing cluster consisting of one or more compute servers dedicated to the production of its products.

4.2.1 Physical view

Physically each cluster is separate. They process data specific to an instrument. Instrument specific Product Generation Executives (PGEs) run on the processors comprising a cluster assigned to an instrument. Each cluster is self-contained and contains all the infrastructure necessary to process a particular instrument. Figure 4.2-1 illustrates the physical view of the topology.

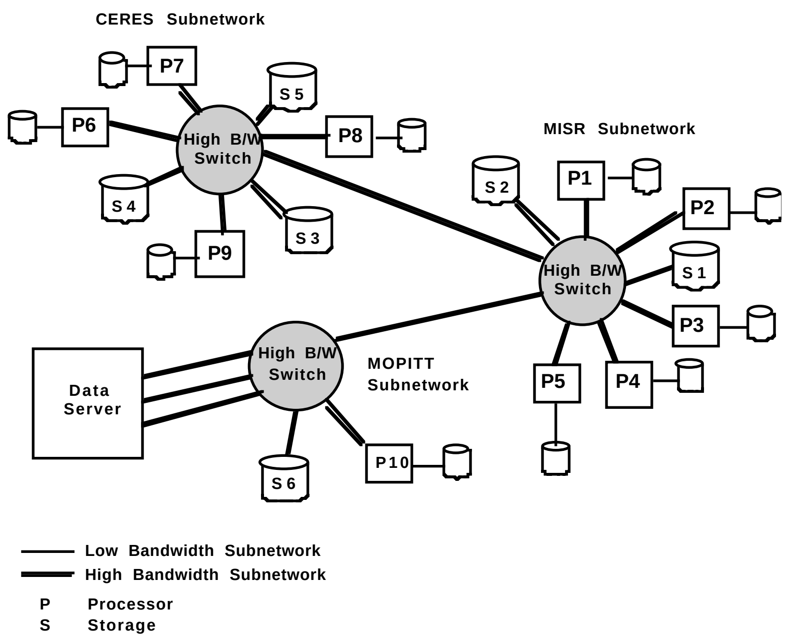

Figure 4.2-1 Physical View of One Instrument's Products Per Cluster

4.3 One instrument's products per cluster except for selected products requiring major processing resources

This option is identical to the first option except that at certain times when processing resources for certain products of that instrument exceed the maximum allowable resources of that processing cluster, then processing of that product can shift to a cluster that has the available resources to support it.

4.3.1 Physical view

A good example of this situation is pooling resources for MISR and CERES Subsystems 4 and 5 (largest subsystems in terms of processing requirements). Figure 4.3-1 shows a high performance subnetwork consisting of high performance machines exclusively for MISR and CERES Subsystems 4 and 5. This subnetwork shares two instruments. Since CERES Subsystems 4 and 5 require data generated by other CERES subsystems (product chains) on the CERES subnetwork, dependent CERES data need to be moved to the high performance subnetwork. Product chains are the concept of including dependent production of data products at higher levels. For example, it is the production of L1 data from L0 data using a L1 algorithm, L2 from L1 and so on. The production of some data products is dependent on other previous level data products and ancillary data products. It should be noted that although the subnetworks are physically separate, there is no logical separation.

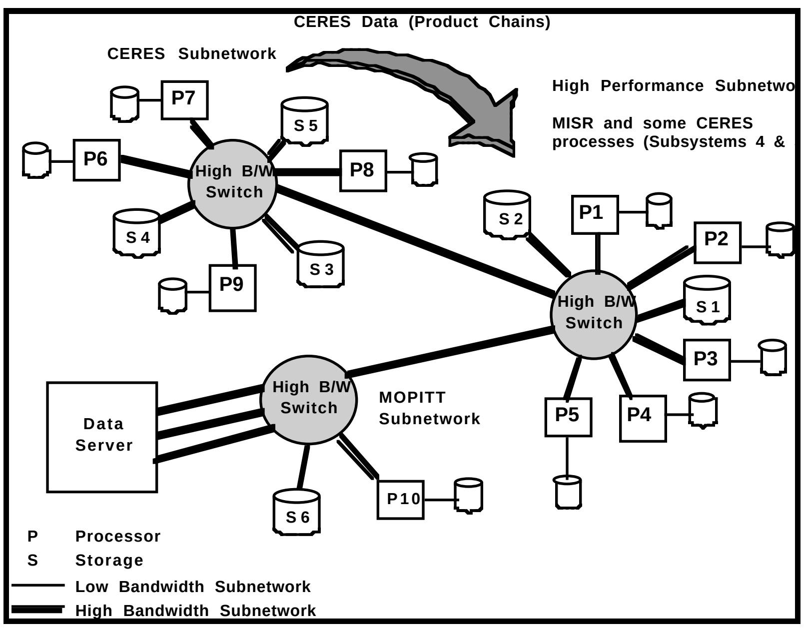

Figure 4.3-1 Physical View of One Instrument's Products Per Cluster Except For Selected Products Requiring Major Processing Resources

4.4 Multiple instruments' products on any cluster

This may apply to conditions whereby instruments with interdependent processing may be collocated. This situation does not apply for the LaRC scenario. There is no interdependency among products from the three instruments. Also, the MODIS products that each instrument uses are different.

4.4.1 Physical view

In the topology presented in Figure 4.4-1, each science processing network has sufficient resources to support multiple instruments. Data need to be moved from one subnetwork to another if there are product chain dependencies. Again, there is no logical separation of the subnetworks. They are only physically separated.

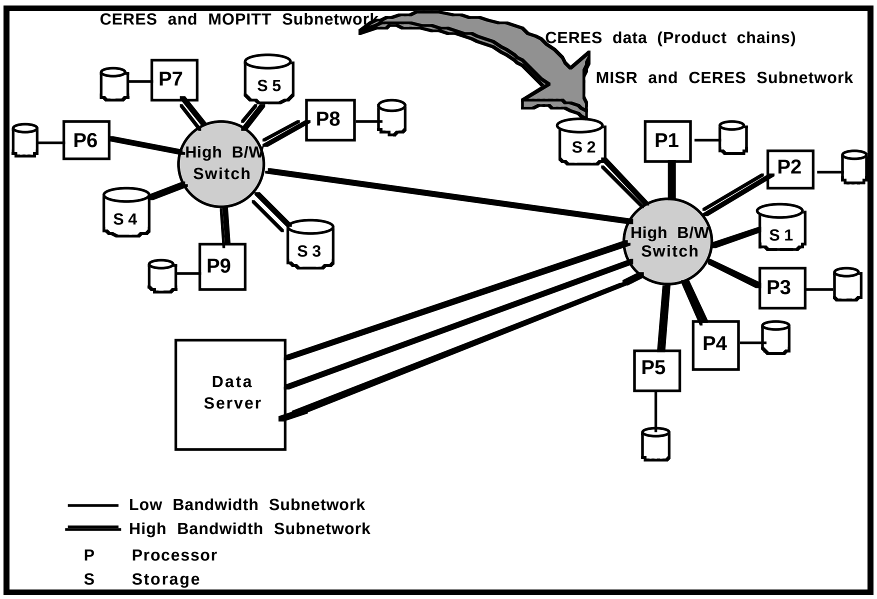

Figure 4.4-1 Physical View of Multiple Instruments' Products On Any Cluster

4.5 Any instruments products on any cluster that can support it

This option is a mix-and-match situation. The processing load will determine the cluster where a particular instrument's products will be processed.

4.5.1 Physical view

Figure 4.5-1 shows a topology where subnetworks may have resources to support any instrument.

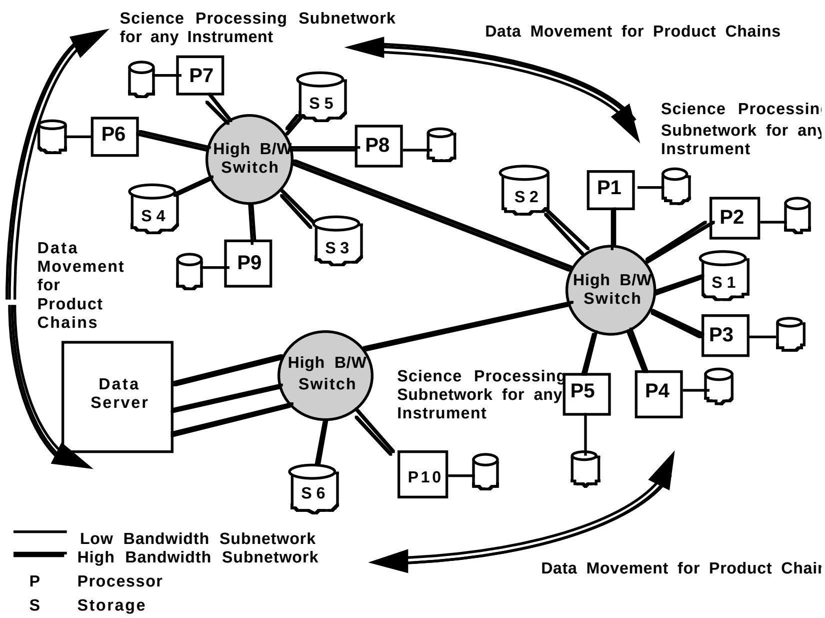

Figure 4.5-1 Physical View Of Any Instruments On Any Cluster That Can Support It

5. Processing Requirements

5.1 Introduction

The January 1995 baseline of the AHWGP data provided the processing requirements for each DAAC site. The data corresponding to epoch "e" (1Q98) represented the Release A requirements, while epoch "k" (3Q99) represented Release B. The following sections give a "rolled up" summary of requirements for each instrument for both epochs. Although, the processing requirements are small at Release A and do not warrant a study of production topologies, nonetheless an analysis is presented here to provide a context.

At Release A, only LaRC and MSFC are production DAACs serving the TRMM mission. Although Release B has many more DAACs, only GSFC and EDC are considered for this study because of their large processing requirements.

5.2 Release A

5.2.1 Processing Requirements for LaRC at Release A

Table 5.2-1 is a "rolled up" summary table of CERES processing requirements at LaRC. Table 5.2-2 lists a summary of definitions for the terms used in Table 5.2-1. A part of the CERES Subsystems 4, 5, 6 and 9 are activated 24.8 times a day, while a part of the subsystems 1, 2 and 12 are activated once a day. Other subsystems processes are also activated once a week. Table 5.2-1 Part 2 lists the daily average I/O and CPU requirements for the LaRC DAAC at Release A (epoch "e"). The I/O requirements/day for CERES is obtained by multiplying the I/O requirement for each process by the number of activations per day summed over all processes. The CPU requirements/day for CERES is obtained by multiplying the CPU requirement for each process by the number of activations per day summed over all processes.

Table 5.2-1. CERES Processing Requirements Summary for Release A (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	Total I/O Req'ments (MB)	CPU Req'ments (MFPOs)	Number of input files	Number of output files	Activations (per day)
CERES 1aT	137.6	87.1	852.0	714.4	852.0	20,790	4	25	1.00
CERES 2aT	314.6	313.6	645.9	331.3	535.8	3,780	3	2	1.00
CERES 2bT	113.5	113.5	227.0	113.5	116.9	1	1	1	0.03
CERES 3aT	114.5	113.5	786.0	671.5	786.0	47,250	3	3	0.03
CERES 4aF	298.2	297.2	543.3	245.1	455.8	34,020	8	2	24.80
CERES 5aF	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	24.80
CERES 5aV	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	4.00
CERES 6aT	221.5	220.5	225.7	4.2	225.7	4,914	3	1	24.80
CERES 6cT	4.2	283.7	8.4	8.4	8.4	4	1	1	0.03
CERES 7aT	2,118.6	2,117.6	3,434.6	1,316.0	3,434.6	680,400	330	40	0.20
CERES 8aT	7,897.0	7,896.0	8,263.4	366.4	8,263.4	226,800	242	2	0.03
CERES 9aTF	155.0	154.0	157.1	2.1	157.1	4,914	3	1	24.80
CERES 9bTF	2.1	2.1	4.2	2.1	4.2	4	1	1	0.03
CERES 10aT	10,317.8	10,316.8	10,882.2	564.4	10,882.2	245,700	1,738	1	0.03
CERES 11aT	91.1	91.1	182.2	91.1	182.2	37,800	1	1	0.10
CERES 12aF	32.9	32.4	284.9	252.0	284.9	37,800	10	24	1.00

Table 5.2-1. CERES Processing Requirements Summary for Release A (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O operations (MB/day)
CERES 1aT	87	714	2.08E+04	852
CERES 2aT	314	331	3.78E+03	536
CERES 2bT	3	3	3.00E-02	4
CERES 3aT	3	20	1.42E+03	24
CERES 4aF	7,371	6,078	8.44E+05	11,305
CERES 5aF	3,819	5,468	6.63E+07	9,312
CERES 5aV	616	882	1.07E+07	1,502
CERES 6aT	5,468	104	1.22E+05	5,597
CERES 6cT	9	0	1.20E-01	0
CERES 7aT	424	263	1.36E+05	687
CERES 8aT	237	11	6.80E+03	248
CERES 9aTF	3,819	52	1.22E+05	3,896
CERES 9bTF	0	0	1.20E-01	0
CERES 10aT	310	17	7.37E+03	326
CERES 11aT	9	9	3.78E+03	18
CERES 12aF	32	252	3.78E+04	285
Total	**22,521**	**14,207**	**7.83E7**	**36,874**

Table 5.2-2. Summary of Definitions

Term	Definition
Volume at initiation	Data volume expected at the start of a process initiation prior to each activation
Volume at completion	Data volume expected at the completion of a process after each activation. It is equal to volume at initiation + sum of all output file sizes for each activation. It includes temporary files.
Staging I/O	Volume staged from the archive for each activation of a process
Destaging I/O	Volume destaged to the archive after completion of a process (for each activation).
I/O requirements	Read and write operations from the staging disk per process for each activation.
CPU requirements	Millions of floating point operations per process per activation. Note that there is no time involved.
Number of input files	Number of input files each process requires per activation. This includes temporary files, input files from other instruments, or other lower level product files from the same instrument.
Number of output files	Number of output files each process produces per activation. This includes temporary files. Some output files may be input to other higher level processes.
Number of activations per day	The number of times the process is activated per day. If the number of activations is a fraction, then it is not activated every day. The process may be activated once every week. To get the number of times it is activated in a month, multiply the number of daily activations by 30.
Volume staged per day	The volume staged for each process times the number of activations of the process per day. This gives a daily average of the volume staged to a process.
Volume destaged per day	The volume destaged for each process times the number of activations of the process per day. This gives a daily average of the volume destaged after completion of a process.
CPU requirements per day	This is a daily average CPU requirements per process. It is obtained by multiplying CPU requirements for each process by the number of activations per day.

5.2.2 Processing Requirements for MSFC at Release A

The MSFC DAAC handles LIS processing at Release A. Table 5.2-3 provides a summary of the LIS processing requirements. LIS processing requirements are small relative to CERES processing.

Table 5.2-3. LIS Processing Requirements Summary for Release A (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	Total I/O Req'ments (MB)	CPU Req'ments (MFPOs)	Number of input files	Number of output files	Activations (per day)
LIS	5.9	5.9	92.2	86.3	91.0	2,492	2	11	14.56

Table 5.2-3. LIS Processing Requirements Summary for Release A (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O operations (MB/day)
LIS	86	1,256	3.63E+04	1,325

5.3 Release B

5.3.1 Processing Requirements for LaRC at Release B

5.3.1.1 ACRIM

The ACRIM processing requirements are not clearly defined in the January 1995 AHWGP baseline. Nonetheless, as indicated in Table 5.3-1, the processing requirements are small.

Table 5.3-1. ACRIM Processing Requirements Summary

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
ACRIM 1A			3.0	3.0	3.0		1	1	15.00

5.3.1.2 CERES

The CERES data processing is organized into ten subsystems. These Subsystems are a logical collection of algorithms which together convert input data products into output data products. Table 5.3-2 lists CERES requirements per process with Table 5.2-2 providing definitions to the terms. Table 5.3-2 and following tables representing the requirements summary are "rolled up" based on the numbers given by the AHWGP. It includes multiple instances of similar files. It also takes into account fractions of files read. As an example, for CERES there can be 240 instances of a file per activation. Subsystems 4 and 5 require the most data volumes at initiation and completion. They also have a substantial Millions of Floating Point Operations (MFPOs). In later sections we will see how these data volumes have implications for external staging disk capacity.

Table 5.3-2. CERES Processing Requirements Summary for Release B (Part 1 of 4)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Outpu t Files	Activations (per day)
CERES 1aA	137.6	87.1	852.0	714.4	852.0	20,790	4	25	1.00
CERES 1aT	137.6	87.1	852.0	714.4	852.0	20,790	4	25	1.00
CERES 1bA	137.6	87.1	852.0	714.4	852.0	20,790	4	25	1.00
CERES 2aA	314.6	313.6	645.9	331.3	535.8	3,780	3	2	1.00
CERES 2aT	314.6	313.6	645.9	331.3	535.8	3,780	3	2	1.00
CERES 2bA	113.5	113.5	227.0	113.5	116.9	1	1	1	0.03
CERES 2bT	113.5	113.5	227.0	113.5	116.9	1	1	1	0.03
CERES 3aA	114.5	113.5	786.0	671.5	786.0	47,250	3	3	0.03
CERES 3aT	114.5	113.5	786.0	671.5	786.0	47,250	3	3	0.03
CERES 3bTA	228.0	227.0	899.5	671.5	899.5	94,500	4	3	0.03
CERES 4aF	298.2	297.2	543.3	245.1	455.8	34,020	8	2	24.80
CERES 4bAF	2,734.8	2,733.8	3,133.9	399.1	3,133.9	3,402,000	9	3	24.80
CERES 5aF	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	24.80
CERES 5aV	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	4.00
CERES 5cAF	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	24.80
CERES 5cAV	155.0	154.0	375.5	220.5	375.5	2,672,460	3	1	4.00
CERES 6aA	221.5	220.5	225.7	4.2	225.7	4,914	3	1	24.80
CERES 6aT	221.5	220.5	225.7	4.2	225.7	4,914	3	1	24.80
CERES 6cA	4.2	292.1	8.4	16.8	8.4	4	1	1	0.03
CERES 6cT	4.2	283.7	8.4	8.4	8.4	4	1	1	0.03
CERES 7aA	2,118.6	2,117.6	3,434.6	1,316.0	3,434.6	680,400	330	40	0.20
CERES 7aT	2,118.6	2,117.6	3,434.6	1,316.0	3,434.6	680,400	330	40	0.20
CERES 7c	2,824.2	2,823.2	4,140.2	1,316.0	4,140.2	1,360,800	498	40	0.03
CERES 8aA	7,897.0	7,896.0	8,263.4	366.4	8,263.4	226,800	242	2	0.03
CERES 8aT	7,897.0	7,896.0	8,263.4	366.4	8,263.4	226,800	242	2	0.03
CERES 8c	7,897.0	7,896.0	8,263.4	366.4	8,263.4	453,600	242	2	0.03

CERES 9aAF	155.0	154.0	157.1	2.1	157.1	4,914	3	1	24.80

Table 5.3-2. CERES Processing Requirements Summary for Release B (Part 2 of 4)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
CERES 9aTF	155.0	154.0	157.1	2.1	157.1	4,914	3	1	24.80
CERES 9bAF	2.1	2.1	4.2	2.1	4.2	4	1	1	0.03
CERES 9bTF	2.1	2.1	4.2	2.1	4.2	4	1	1	0.03
CERES 10aA	10,317.8	10,316.8	10,882.2	564.4	10,882.2	245,700	1,738	1	0.03
CERES 10aT	10,317.8	10,316.8	10,882.2	564.4	10,882.2	245,700	1,738	1	0.03
CERES 10bTA	11,880.2	11,879.2	12,444.6	564.4	12,444.6	491,400	2,482	1	0.03
CERES 11aA	91.1	91.1	182.2	91.1	182.2	37,800	1	1	0.10
CERES 11aT	91.1	91.1	182.2	91.1	182.2	37,800	1	1	0.10
CERES 12aF	32.9	32.4	284.9	252.0	284.9	37,800	10	24	1.00

Table 5.3-2. CERES Processing Requirements Summary for Release B (Part 3 of 4)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O operations (MB/day)
CERES 1aA	87	714	2.08E+04	852
CERES 1aT	87	714	2.08E+04	852
CERES 1bA	87	714	2.08E+04	852
CERES 2aA	314	331	3.78E+03	536
CERES 2aT	314	331	3.78E+03	536
CERES 2bA	3	3	3.00E-02	4
CERES 2bT	3	3	3.00E-02	4
CERES 3aA	3	20	1.42E+03	24
CERES 3aT	3	20	1.42E+03	24
CERES 3bTA	7	20	2.84E+03	27
CERES 4aF	7,371	6,078	8.44E+05	11,305
CERES 4bAF	67,798	9,898	8.44E+07	77,721

CERES 5aF	3,819	5,468	6.63E+07	9,312
CERES 5aV	616	882	1.07E+07	1,502

Table 5.3-2. CERES Processing Requirements Summary for Release B (Part 4 of 4)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O operations (MB/day)
CERES 5cAF	3,819	5,468	6.63E+07	9,312
CERES 5cAV	616	882	1.07E+07	1,502
CERES 6aA	5,468	104	1.22E+05	5,597
CERES 6aT	5,468	104	1.22E+05	5,597
CERES 6cA	9	1	1.20E-01	0
CERES 6cT	9	0	1.20E-01	0
CERES 7aA	424	263	1.36E+05	687
CERES 7aT	424	263	1.36E+05	687
CERES 7c	85	39	4.08E+04	124
CERES 8aA	237	11	6.80E+03	248
CERES 8aT	237	11	6.80E+03	248
CERES 8c	237	11	1.36E+04	248
CERES 9aAF	3,819	52	1.22E+05	3,896
CERES 9aTF	3,819	52	1.22E+05	3,896
CERES 9bAF	0	0	1.20E-01	0
CERES 9bTF	0	0	1.20E-01	0
CERES 10aA	310	17	7.37E+03	326
CERES 10aT	310	17	7.37E+03	326
CERES 10bTA	356	17	1.47E+04	373
CERES 11aA	9	9	3.78E+03	18
CERES 11aT	9	9	3.78E+03	18
CERES 12aF	32	252	3.78E+04	285
Total	**106,209**	**32,783**	**2.4E+08**	**136,940**

5.3.1.3 MISR

The MISR processing requirements summary is listed in Table 5.3-3 (the terms in each column are defined in Table 5.2-2). MISR data will be processed in units of one orbit. This translates to approximately 14.56 activations per day for all MISR processes. There will be four production software subsystems, one each for the products at Levels 1A, 1b, 2-T/C, and 2-A/S. Each of these will be capable of being operated individually, or as a combined unit that maximizes resources and throughput. Processing will not commence until the Planning Subsystem determines that all the data dependencies are satisfied. Operational data from external resources, e.g., meteorological data from NOAA or instrument data from MODIS will be preprocessed by a separate element of the respective MISR software subsystems to prepare this data for use.

Table 5.3-3. MISR Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
MISP1A	3,167.2	3,167.2	7,451.2	4,284.1	7,451.2	237,000	23	76	14.50
MISP1B1	3,815.8	3,815.8	7,745.0	3,929.2	7,745.0	178,000	39	38	14.50
MISP1B2	7,624.0	7,624.0	11,989.0	4,365.0	11,989.0	8,557,000	59	22	14.50
MISP2AS	5,973.6	5,962.7	6,455.4	481.8	6,455.4	6,990,000	48	4	14.50
MISP2TC	4,449.7	4,449.7	4,725.2	275.5	4,725.2	4,652,000	40	2	14.50

Table 5.3-3. MISR Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O operations (MB/day)
MISP1A	45,924	62,119	3.44E+06	108,043
MISP1B1	55,329	56,974	2.58E+06	112,303
MISP1B2	110,549	63,293	1.24E+08	173,841
MISP2AS	86,459	6,986	1.01E+08	93,603
MISP2TC	64,521	3,995	6.75E+07	68,515
Total	**362,781**	**193,365**	**2.99E8**	**556,305**

5.3.1.4 MOPITT

The scientific goals for MOPITT depend upon long term, homogeneous, global data products rather than the quick turn around of observations made at a particular geographical location. Therefore, the MOPITT Standard Product Generation algorithms will be designed in such a way that each level of processing is autonomous. Each level is dependent on the existence of its preceding level and upon the existence of certain ancillary data unique to a level. Processors for the various levels may be run in sequence or at different times assuming that the necessary dependencies are met. Table 5.3-4 lists the summary of requirements for generating MOPITT standard products. The terms in each column of Table 5.4-4 are defined in Table 5.2-2. It is expected that L1 (Calibrated, Earth Located Radiance) and L2 (Retrieved Geophysical Parameters) data products will be produced on a daily basis (see activations per day column in Table 5.3-4). The minimum volume of data to be handled for each run will be data accumulated within the corresponding day. The L3 (Global, Gridded, Geophysical Parameters) products will be produced on a weekly basis.

Table 5.3-4. MOPITT Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
MOPL1	257.2	255.2	359.2	101.0	359.2	16,800	3	3	1.00
MOPL1Qi-D	101.0	101.0	111.0	10.0	111.0	900	2	1	1.00
MOPL2-E	31,571.6	31,311.4	31,756.3	184.7	31,756.3	1,502,250	1,186	3	1.00
MOPL2Qi-D	174.7	174.7	184.7	10.0	184.7	1,350	2	1	1.00
MOPL3	522.9	522.9	611.5	88.6	611.5	28,290	7	3	0.14
MOPL3Qi-F	78.6	78.6	88.6	10.0	88.6	900	2	1	0.14

Table 5.3-4. MOPITT Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
MOPL1	255	101	1.68E+04	359
MOPL1Qi-D	101	10	9.00E+02	111
MOPL2-E	31,311	185	1.50E+06	31,756
MOPL2Qi-D	175	10	1.35E+03	185
MOPL3	74	13	4.02E+03	87

MOPL3Qi-F	11	1	1.28E+02	13
Total	**31,928**	**320**	**1.53E6**	**32,511**

5.3.1.5 Total Processing Requirements at LaRC

The LaRC DAAC may contain three processing clusters with each one corresponding to one instrument. The actual configuration may vary. MODIS products from EDC and GSFC are used as ancillary inputs. Tables 5.3-2 to 5.3-4 (Part 2) lists the daily average I/O and CPU requirements at the LaRC DAAC for each instrument (CERES, MISR and MOPITT) at epoch "k". The I/O requirements/day for CERES is obtained by multiplying the I/O requirement for each process by the number of activations per day summed over all processes. The CPU requirements/day for CERES is obtained by multiplying the CPU requirement for each process by the number of activations per day summed over all processes. Similarly, the daily requirements for MISR and MOPITT are determined.

5.3.2 Processing requirements by instrument at GSFC

5.3.2.1 COLOR

The COLOR 1B and the Global Area Coverage (GAC) processes are activated for every orbit of data (14.56 times a day) as indicated in Table 5.3-5. All other processes are activated only once a day.

Table 5.3-5. COLOR Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
CLR Aero Browse	720.0	720.0	721.0	1.0	721.0	50	15	1	1.00
CLR Chl. Browse	720.0	720.0	721.0	1.0	721.0	50	15	1	1.00
CLR K 490 Browse	720.0	720.0	721.0	1.0	721.0	50	15	1	1.00
CLR Rad. Browse	720.0	720.0	729.0	9.0	729.0	50	15	1	1.00
CLR WLR Browse	720.0	720.0	721.0	1.0	721.0	50	15	1	1.00
COLOR 1B	42.0	42.0	90.0	48.0	90.0	123,624	1	1	14.56
COLOR Aerosol	720.0	720.0	757.0	37.0	757.0	43,632	15	1	1.00
COLOR Ancillary	720.0	720.0	721.0	1.0	721.0	50	15	1	1.00
COLOR Binned	720.0	720.0	1,107.0	387.0	1,107.0	453,686	15	1	1.00
COLOR Chloro-A	720.0	720.0	757.0	37.0	757.0	43,632	15	1	1.00

COLOR GAC	48.0	48.0	80.0	32.0	80.0	44,565	1	1	14.56
COLOR K 490	720.0	720.0	757.0	37.0	757.0	43,632	15	1	1.00
COLOR W.L. Rad.	720.0	720.0	757.0	37.0	757.0	43,632	15	1	1.00

Table 5.3-5. COLOR Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
CLR Aero Browse	720	1	5.00E+01	721
CLR Chl. Browse	720	1	5.00E+01	721
CLR K 490 Browse	720	1	5.00E+01	721
CLR Rad. Browse	720	9	5.00E+01	729
CLR WLR Browse	720	1	5.00E+01	721
COLOR 1B	612	699	1.80E+06	1,310
COLOR Aerosol	720	37	4.36E+04	757
COLOR Ancillary	720	1	5.00E+01	721
COLOR Binned	720	387	4.54E+05	1,107
COLOR Chloro-A	720	37	4.36E+04	757
COLOR GAC	699	466	6.49E+05	1,165
COLOR K 490	720	37	4.36E+04	757
COLOR W.L. Rad.	720	37	4.36E+04	757
Total	**9,230**	**1,714**	**3.08E6**	**10,944**

5.3.3.2 MODIS

Table 5.3-6 lists the MODIS processing requirements. The MODIS L1 processes (1A production from L0 and 1B production of calibrated radiances) are activated 585 times a day. The 1B production of calibrated radiances is CPU intensive. The MODIS L2 production of cloud product (cloud optical and cloud top properties) are activated 586 times a day. The MODIS L3 compositing of daily, weekly and monthly aerosol products using the L3 processes are activated up to 51 times a day. The dynamic analysis does not simulate MODIS L3 processes because of uncertainties in data required by the dynamic model. Clearly the L3 processes require not only large staging/destaging volumes, but also CPU requirements. The MODIS Ocean Color production of daily and weekly composites require large staging/destaging volumes and I/O.

Table 5.3-6. MODIS (GSFC) Processing Requirements Summary (Part 1 of 4)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
MOD01:L1A:G	120.0	120.0	316.8	196.8	316.8	14,757	1	1	585.00
MOD02:L1B:G	196.8	196.8	504.7	307.9	504.7	191,836	1	1	585.00
MOD03:L1A:G	304.6	196.8	328.6	23.0	45.6	6,050	4	2	585.00
MOD04:L2:G	376.6	374.6	380.4	3.7	380.4	184	7	1	292.50
MOD04:L3:DY:G	41.0	41.0	44.5	3.5	44.5	39	11	1	50.71
MOD04:L3:MN:G	13.8	13.8	17.3	3.5	17.3	11	4	1	11.83
MOD04:L3:WK:G	26.1	26.1	29.6	3.5	29.6	19	7	1	50.71
MOD05:L2:G	342.4	336.4	380.3	37.9	380.3	28	9	1	292.50
MOD06:L2:G	966.9	498.1	984.4	17.5	556.9	10,705	19	2	585.00
MOD06:L3:MN:G	5,396.2	5,396.2	5,396.2	0.0	5,396.2	110	308	1	11.83
MOD09:L2:I	544.5	344.5	685.5	141.0	495.5	2,286	5	1	292.50
MOD10:L2:I	408.6	208.6	410.0	1.4	220.0	53	6	1	292.50
MOD11:L2:I	805.6	404.6	816.4	10.8	437.5	640	8	1	585.00
MOD11:L3:WK:G	820.8	820.8	832.8	12.0	832.8	147	76	1	50.71
MOD13:L2:G	658.1	458.1	712.1	54.0	522.1	1,914	5	1	292.50
MOD14:L2:G	586.1	586.1	588.8	2.7	588.8	24	8	1	585.00
MOD28:D:ORBIT:G	27.1	27.1	81.8	54.7	81.8	4,050	20	1	15.00
MOD28:L2:G	337.7	337.7	343.1	5.4	343.1	30,465	5	1	585.00
MOD28:L3:COMP:D:DY:G	874.7	874.7	1,271.1	396.4	1,271.1	8,100	16	1	1.00
MOD28:L3:COMP:N:DY:G	874.7	874.7	1,271.1	396.4	1,271.1	8,100	16	1	1.00
MOD28:L3:D:WK:G	3,171.0	3,171.0	6,342.1	3,171.0	6,342.1	27,000	8	8	0.14
MOD28:L3:N:WK:G	3,171.0	3,171.0	6,342.1	3,171.0	6,342.1	27,000	8	8	0.14
MOD28:L3:TMP:D:WK:G	3,567.4	3,567.4	4,360.2	792.8	4,360.2	13,500	9	2	0.14

Table 5.3-6. MODIS (GSFC) Processing Requirements Summary (Part 2 of 4)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
MOD28:L3:TMP:N:WK:G	3,567.4	3,567.4	4,360.2	792.8	4,360.2	27,000	9	2	0.29
MOD28:N:ORBIT:G	27.1	27.1	81.8	54.7	81.8	4,050	20	1	15.00
MOD28:SPBIN:G	5.4	5.4	6.8	1.4	6.8	10,815	1	1	585.00
MOD29:L2:G	162.6	162.6	165.3	2.7	165.3	24	4	1	292.50
MOD32:L2:G	491.7	486.7	497.7	6.0	497.7	1	9	0	1.00
MOD35:L2:G	726.9	356.9	732.4	5.4	494.9	8,952	10	1	585.00
MOD41:L2:H	48.6	47.6	129.6	81.0	129.6	789	5	1	292.50
MOD:ATMOS:L2:G	796.9	346.9	834.2	37.3	406.7	1,476	7	4	585.00
MOD:ATMOS:L3:MN:G	458.9	458.9	458.9	0.0	458.9	74	924	1	11.83
MODOCCLR:L2:G	329.8	329.3	497.7	167.9	497.7	82,676	6	1	292.50
MODOCCLR:L3:COMP:DY:G	13,993.6	13,993.6	26,675.6	12,682.0	26,675.6	12,384	16	1	1.00
MODOCCLR:L3:TMP:WK:G	114,138.0	102,767.6	50,728.0	25,364.0	50,728.0	20,640	9	2	0.14
MODOCCLR:L3:WK:QC:G	101,456.0	101,456.0	202,912.0	101,456.0	202,912.0	41,280	8	8	0.14
MODOCCLR:ORBIT:G	433.2	433.2	1,307.8	874.6	1,307.8	6,192	20	1	14.00
MODOCCLR:SPBIN:G	167.9	167.9	189.6	21.7	189.6	16,535	1	1	292.50

Table 5.3-6. MODIS (GSFC) Processing Requirements Summary (Part 3 of 4)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
MOD01:L1A:G	70,200	115,099	8.63E+06	185,299
MOD02:L1B:G	115,099	180,145	1.12E+08	295,244
MOD03:L1A:G	115,099	13,455	3.54E+06	26,651
MOD04:L2:G	109,582	1,091	5.38E+04	111,258
MOD04:L3:DY:G	2,081	175	1.98E+03	2,256
MOD04:L3:MN:G	164	41	1.30E+02	205
MOD04:L3:WK:G	1,324	175	9.64E+02	1,500
MOD05:L2:G	98,385	11,089	8.19E+03	111,229
MOD06:L2:G	291,383	10,249	6.26E+06	325,792
MOD06:L3:MN:G	63,855	0	1.30E+03	63,855
MOD09:L2:I	100,755	41,243	6.69E+05	144,922
MOD10:L2:I	61,021	410	1.54E+04	64,356
MOD11:L2:I	236,662	6,318	3.74E+05	255,908
MOD11:L3:WK:G	41,626	609	7.47E+03	42,235
MOD13:L2:G	133,991	15,795	5.60E+05	152,711
MOD14:L2:G	342,839	1,580	1.39E+04	344,419
MOD28:D:ORBIT:G	406	820	6.08E+04	1,226
MOD28:L2:G	197,572	3,159	1.78E+07	200,731
MOD28:L3:COMP:D:DY:G	875	396	8.10E+03	1,271
MOD28:L3:COMP:N:DY:G	875	396	8.10E+03	1,271
MOD28:L3:D:WK:G	444	444	3.78E+03	888
MOD28:L3:N:WK:G	444	444	3.78E+03	888
MOD28:L3:TMP:D:WK:G	499	111	1.89E+03	610
MOD28:L3:TMP:N:WK:G	1,019	227	7.71E+03	1,246
MOD28:N:ORBIT:G	406	820	6.08E+04	1,226

MOD28:SPBIN:G	3,159	792	6.33E+06	3,951
MOD29:L2:G	47,566	790	7.02E+03	48,356

Table 5.3-6. MODIS (GSFC) Processing Requirements Summary (Part 4 of 4)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
MOD32:L2:G	487	6	1.00E+00	498
MOD35:L2:G	208,810	3,171	5.24E+06	289,493
MOD41:L2:H	13,929	23,693	2.31E+05	37,914
MOD:ATMOS:L2:G	202,960	21,809	8.63E+05	237,931
MOD:ATMOS:L3:MN:G	5,431	0	8.70E+02	5,431
MODOCCLR:L2:G	96,332	49,108	2.42E+07	145,586
MODOCCLR:L3:COMP:DY:G	13,994	12,682	1.24E+04	26,676
MODOCCLR:L3:TMP:WK:G	14,681	3,623	2.95E+03	7,247
MODOCCLR:L3:WK:QC:G	14,494	14,494	5.90E+03	28,987
MODOCCLR:ORBIT:G	6,065	12,244	8.67E+04	18,309
MODOCCLR:SPBIN:G	49,108	6,336	4.84E+06	55,443
Total	**2,663,620**	**553,037**	**1.92E08**	**3,243,020**

5.3.3 Processing requirements by instrument at EDC

5.3.3.1 ASTER

Table 5.3-7 lists ASTER processing requirements. Processes 9 (polar cloud map) and 10 (Digital Elevetion Model -DEM) are CPU intensive while 4 (atmospheric correction - VNIR and SWIR) and 5 (atmospheric correction - TIR) are I/O intensive.

Table 5.3-7. ASTER Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
AST_PGE_01	774.5	243.0	792.5	18.0	342.0	4,993	6	1	182.00
AST_PGE_02	293.0	261.0	360.1	67.1	360.1	18,337	4	3	182.00
AST_PGE_03	243.9	243.0	249.9	6.0	249.6	861	3	1	70.00
AST_PGE_04	11,096.1	764.4	11,571.7	475.6	1,568.8	34,200	13	4	70.00
AST_PGE_05	1,114.1	782.4	1,296.1	182.0	1,293.1	5,221	13	2	70.00
AST_PGE_06	27.6	24.0	35.6	8.0	35.6	3,329	4	1	70.00
AST_PGE_07	327.0	327.0	345.0	18.0	345.0	1,664	4	1	28.00
AST_PGE_08	345.0	327.0	381.0	36.0	381.0	3,329	5	1	28.00
AST_PGE_09	543.2	253.0	561.2	18.0	468.3	102,600	7	1	6.00
AST_PGE_10	243.0	243.0	278.0	35.0	278.0	513,000	1	1	1.00

Table 5.3-7. ASTER Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
AST_PGE_01	44,226	3,276	9.09E+05	62,235
AST_PGE_02	47,502	12,212	3.34E+06	65,538
AST_PGE_03	17,010	420	6.02E+04	17,475
AST_PGE_04	53,507	33,292	2.39E+06	109,813
AST_PGE_05	54,767	12,740	3.65E+05	90,520
AST_PGE_06	1,680	560	2.33E+05	2,492
AST_PGE_07	9,156	504	4.66E+04	9,660
AST_PGE_08	9,156	1,008	9.32E+04	10,668
AST_PGE_09	1,518	108	6.16E+05	2,810
AST_PGE_10	243	35	5.13E+05	278
Total	**238,766**	**64,155**	**8.57E6**	**371,489**

5.3.3.2 MODIS at EDC

As shown in Table 5.3-8, the MODIS L3 production of Land Cover is I/O intensive. The MODIS L3 production of BRDF/Albedo is a very CPU intensive process activated approximately 39 times a day.

Table 5.3-8. MODIS (EDC) Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
MOD09:L3:9DY:G	65,712.5	65,212.5	66,792.5	1,080.0	65,549.5	2,803,741	401	1	39.44
MOD12:L3:3MN:I	487,188.8	486,988.8	487,217.2	28.4	481,715.1	280,338	4,023	1	3.94
MOD14:L3:10DY:G	27.0	27.0	29.7	2.7	29.7	27,042	10	1	35.50
MOD14:L3:DY:G	29.7	29.7	32.4	2.7	32.4	1,242	11	1	355.00
MOD14:L3:MN:G	8.1	8.1	10.8	2.7	10.8	736	3	1	11.83
MOD15:L4:WK:G	4,986.4	4,186.4	5,789.4	3.0	4,205.4	124	79	2	50.71
MOD16:L3:WK:G	16.0	15.0	102.4	86.4	101.4	32,710	3	1	50.71
MOD17:L4:WK:G	49.4	49.4	222.2	172.8	222.2	126	5	1	50.71
MOD34:L3:10DY:I	8,300.0	8,100.0	8,314.4	14.4	8,124.4	27,042	113	1	35.50
MOD34:L3:MN:I	43.2	43.2	129.6	86.4	129.6	3,504	3	1	11.83
MOD40:L3:10DY:I	8.6	8.6	9.5	0.9	9.5	384	10	1	35.50
MOD40:L3:DY:I	1,579.5	1,579.5	1,580.4	0.9	1,580.4	12	585	1	355.00
MOD40:L3:MN:I	2.6	2.6	3.4	0.9	3.4	117	3	1	11.83

Table 5.3-8. MODIS (EDC) Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
MOD09:L3:9DY:G	2,572,271	42,600	1.11E+08	2,585,564
MOD12:L3:3MN:I	1,920,900	112	1.11E+06	1,900,098
MOD14:L3:10DY:G	959	96	9.60E+05	1,054
MOD14:L3:DY:G	10,544	959	4.41E+05	11,502
MOD14:L3:MN:G	96	32	8.71E+03	128
MOD15:L4:WK:G	212,310	152	6.31E+03	213,274
MOD16:L3:WK:G	761	4,382	1.66E+06	5,143
MOD17:L4:WK:G	2,505	8,763	6.41E+03	11,269
MOD34:L3:10DY:I	287,550	511	9.60E+05	288,416
MOD34:L3:MN:I	511	1,022	4.15E+04	1,534
MOD40:L3:10DY:I	305	31	1.36E+04	336
MOD40:L3:DY:I	560,723	305	4.43E+03	561,028
MOD40:L3:MN:I	31	10	1.38E+03	41
Total	**5,569,465**	**58,975**	**1.16E8**	**5,579,386**

5.3.4 Processing Requirements at MSFC

5.3.4.1 LIS

The LIS L1, L2 and L3 processes are activated for every orbit of data (14.56 times a day). LIS processing requirements are small.

Table 5.3-9. LIS Processing Requirements Summary (Part 1 of 2)

Process	Volume at Initiation (MB)	Staging I/O (MB)	Volume at Completion (MB)	Destaging I/O (MB)	I/O Reqts. (MB)	CPU Reqts. (MFPOs)	No. Input Files	No. Output Files	Activations (per day)
LIS	5.9	5.9	92.2	86.3	91.0	2,492	2	11	14.56

Table 5.3-9. LIS Processing Requirements Summary (Part 2 of 2)

Process	Volume Staged (MB/day)	Volume Destaged (MB/day)	CPU requirements (MFPOs)/day	I/O Operations (MB/day)
LIS	86	1,256	3.63E+04	1,325

6. Analysis By Cluster Optimization Alternative

6.1 General Assumptions

The following assumptions are made for the analysis of AHWGP data:

- The AHWGP (January 1995 baseline) data are representative of the kind of processing to be supported. As science algorithms are developed and launch dates near, these data are subject to change;
- Only generation of standard products is considered. Reprocessing will be considered as more AHWGP data become available, and the ECS Systems Performance Model is made more sophisticated;
- A 24-hour time period is assumed for all Release B calculations;
- All estimates for computational load are based on Millions of Floating Point Operations (MFPOs). No distinction is made between floating point operations and non-floating point operations. They are two entirely different machine attributes that can vary within an architecture;
- Only raw numbers from the AHWGP are considered for this study. Input data are not scaled to represent the increase in processing requirements required for hardware selection;
- I/O control delay is not accounted for in static computations of processing times;
- V0 data migration and user pull are not considered for the static analysis.

6.2 Cluster Optimization Alternatives Drivers

The choice of optimization alternative is based on the following drivers:

- Processing requirements (MFLOPS) of various instruments;
- Product chains - production of data products at the next higher level. For example, it is the production of L1 data from L0 data using L1 algorithm, L2 from L1 and so on. The production of some data products is dependent on other previous level data products and ancillary data products;
- Special hardware or software requirements for each instrument;
- Process activations and estimated length of process runs;
- Anticipated growth in individual instrument requirements;
- Interdependency among multiple instruments (one instruments products is input to another instrument).

- Reduction in the complexity of data flow (i.e. minimize data hops as much as possible)

In the following sections, the AHWGP data are analyzed both statically and dynamically. A detailed analysis is performed only for the "one instrument's products per cluster" alternative. The other alternatives will be analyzed when the ECS Systems Performance Model has attained sufficient maturity to handle them.

6.3 Analysis by optimization alternative

6.3.1 One instrument's products per cluster

Processing each instrument on a separate cluster is a natural extension of the processing clusters topology. Unique algorithm requirements may dictate the selection of particular compute resources that offer the best solutions.

6.3.1.1 Static analysis

6.3.1.1.1 Estimating Number of Processors and Staging Disk Requirement

A host attached disk is assumed for the calculations. The theoretical staging disk capacity can be estimated based on the "duty cycle" which can be defined as the amount of MFLOPS to be performed by a processor to the number of MFLOPS which a processor is capable. Processor performance is usually rated by vendors as peak MFLOPS. An efficiency factor ($\eta = 0.25$) is used to adjust the vendor provided peak MFLOPS.

$$\text{Duty cycle} = MFPOs|_{process}\ \text{per day} / (N \times MFLOPS|_{processor} \times \eta \times 86400) \text{....... (1)}$$

where $MFPOs|_{process}$ per day : Millions of Floating Point Operations performed by the process per day,

N: Number of processors

$MFLOPS|_{processor}$: Peak processor rating per processor.

A duty cycle of unity indicates a fully utilized processor.

The staging disk capacity can be written as:

$$\text{Staging disk requirement (per day)} = Vc \times t_{cp} \text{ .... (2)}$$

where Vc : Data volume at completion (MB);

t_{cp} : Estimated time of completion of a process.

t_{cp} : $MFPOs\ \text{per day} / (N \times MFLOPS|_{processor} \times \eta)$ .... (3)
for duty cycle unity

6.3.1.1.2 Needed I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Host Attached Storage - Static Estimates

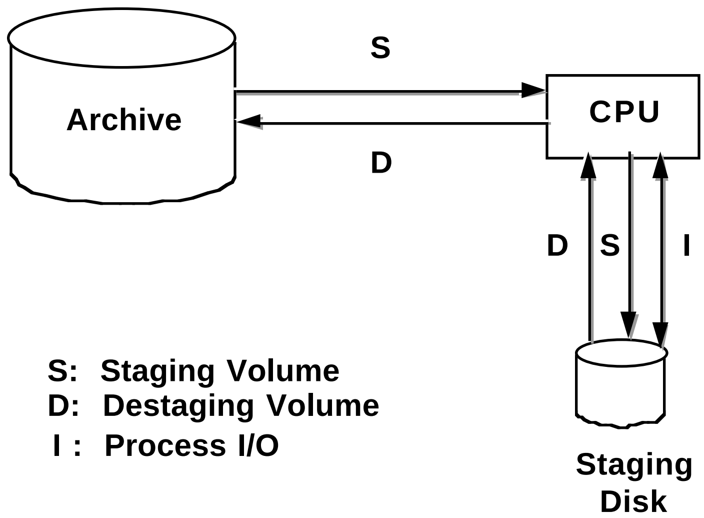

Figure 6.3-1a Schematic of Data Movement at CPU for Host Attached Storage

The I/O bandwidth at CPU gives a fairly good static estimate of the amount of data movement coordinated by a CPU. Figure 6.3-1a is a schematic depicting data movement among the archive, CPU and the staging disk. A host attached staging disk is assumed. When the Planning Subsystem determines that all dependencies are satisfied, it sends a Data Processing Request to the Data Processing Subsystem. Data are then staged to the staging disk by the processing CPU. After processing is completed, the output files that are to be archived are destaged from the staging disk. During processing the CPU coordinates the read/write operations of the process. Therefore, it is possible to theoretically estimate a daily average I/O bandwidth at CPU for each process. The I/O bandwidth at CPU can be formulated as:

$$I/O\ bandwidth\ at\ CPU = (2 \times V_S) + (2 \times V_{ds}) + V_{I/O} \qquad (4)$$

Where V_S: Volume staged;

V_{ds}: Volume destaged;

$V_{I/O}$: Volume of I/O operations.

The data movement between the Data Handler (archive) and CPU gives the theoretical Processor <--> Data Handler throughput, which can be estimated by:

$$Theoretical\ Processor <\!\!-\!\!-\!\!> Data\ Handler\ throughput = V_S + V_{ds} \qquad (5)$$

6.3.1.1.3 Needed I/O Bandwidth at CPU for Network Attached Storage - Static Estimates

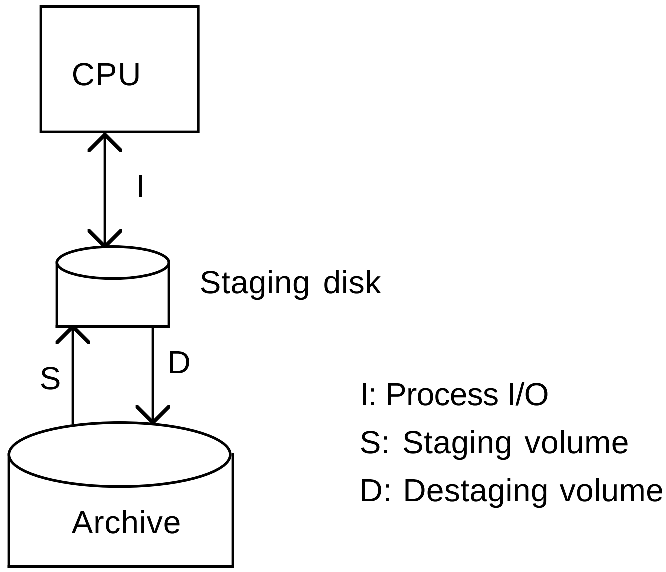

Figure 6.3-1b. Schematic of Data Movement at CPU for Network Attached Storage

The I/O bandwidth at CPU gives a fairly good static estimate of the amount of data movement coordinated by a CPU. Figure 6.3-1b is a schematic depicting data movement among the archive, CPU and the staging disk which is a Network Attached Storage device currently under consideration for Release B. Data movement can be significantly different for Network Attached Storage [4] [5]. When the Planning Subsystem determines that all dependencies are satisfied, it sends a Data Processing Request to the Data Processing Subsystem. Data are then staged to the staging disk. In a network attached device, the science processor does not coordinate staging and destaging. The only contribution to the needed bandwidth at CPU is the due to the process I/O to the staging disk. It is possible to theoretically estimate a daily average I/O bandwidth at CPU for each process. The I/O bandwidth at CPU for Network Attached Storage can be formulated as:

$$I/O \text{ bandwidth at CPU} = V_{I/O} \quad \text{...}(6)$$

Where $V_{I/O}$: Volume of I/O operations.

6.3.1.1.4 Static Analysis by Instrument

6.3.1.1.4.1 CERES

The relationship between duty cycle and the number of processors (for different peak MFLOP ratings) is illustrated in Table 6.3-1a for CERES based on total CPU requirements (see Table 5.2-1 Part 2 converted to per sec). At 300-MFLOP peak rating, we need 12.1 processors at duty cycle 1 (ideally) to satisfy CERES CPU requirements. To be more realistic, this is equivalent to

15.1 processors at 80% duty cycle. Similarly, 6.0 processors with 600-MFLOPS peak rating at duty cycle 1 (ideal case) is equivalent to 7.5 processors at 80% duty cycle. With 900-MFLOPS peak processors, we need only 1.0 additional processor (the difference of 5.0 and 4.0) to allow for a 20% slack in duty cycle.

Table 6.3-1a. CERES Daily Average Processor Requirements as a Function of Duty Cycle for Release A

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	12.1	6.0	4.0
0.95	12.7	6.3	4.2
0.90	13.4	6.7	4.4
0.85	14.2	7.1	4.7
0.80	15.1	7.5	5.0

Similarly, the daily average processor requirements for CERES as a function of duty cycle is given in Table 6.3-1b.

Table 6.3-1b. CERES Daily Average Processor Requirements as a Function of Duty Cycle for Release B

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	37.1	18.5	12.4
0.95	39.1	19.5	13.1
0.90	41.2	20.6	13.8
0.85	43.6	21.8	14.6
0.80	46.4	23.1	15.5

If the duty cycle of a group of processors is held constant at unity, from equation 1 we see that by increasing the number of processors we can process more floating point operations per day. The capacity to process more floating point operations per day is related to throughput which in turn has implications to the staging disk requirement. If throughput can be increased, we can reduce the staging disk requirement. Let us do some theoretical calculations based on the number of processors and determine staging disk requirement given in equation 2.

For a duty cycle of unity, to compute the daily average staging disk requirement we multiply the volume at completion for each CERES process by the number of activations per day (in Tables 5.2-1 Part 2 for Release A and 5.3-2 for Release B) and sum over all processes. The daily average staging disk requirement determined by this method yields 139,330 MB for CERES at Release B. This staging disk capacity is needed when 37.1 300-MFLOP (peak) processors are used at duty cycle unity. The theoretical estimate of staging disk requirement also indicates that 37.1 300-MFLOPS (peak) and 12.4 900-MFLOPS (peak) processors are equivalent. This is true only in the static sense because dynamically, processes on a 900-MFLOPS processor complete faster than on 300-MFLOPS processors, thereby requiring shorter residency times in the disk. The static estimate of staging disk requirement may be an overestimate because it assumes that all processes start at the same time without accounting for disk volumes cleared by processes that end at various times. Nonetheless, it gives an upper bound for validating the dynamic model.

It is possible to determine a relationship between the number of processors and staging disk volume. From Tables 5.2-1 and 5.3-2, based on the number of CPU requirements per day for CERES, we can calculate the time required for completion of all processes. If we use 47 300-MFLOP processors (at duty cycle 1) instead of 37.1, we can increase the processing capacity and reduce the time it takes for all CERES processes to complete. Therefore, the time of completion can be written as:

$$\text{Time of completion} = \text{MFPOs per day} / (N \times \text{MFLOPS}|_{processor} \times) \text{ ....(3);}$$

$$\text{for Release A} = 7.83 \times 10^{7} / (15.0 \times 300.0 \times 0.25 \times 3600 \times 24)$$

$$= 0.805 \text{ days}$$

$$\text{for Release B} = 2.4 \times 10^{8} / (47.0 \times 300.0 \times 0.25 \times 3600 \times 24)$$

$$= 0.788 \text{ days}$$

Similarly, substituting the appropriate numbers into equation (2):

$$\text{Staging disk requirement for Release A (push only)} = 36{,}874 \text{ MB per day} \times 0.805 \text{ days}$$

$$= 29{,}683 \text{ MB} \equiv 30 \text{ GB}$$

$$\text{Staging disk requirement for Release B (push only)} = 139{,}336 \text{ MB per day} \times 0.788 \text{ days}$$

$$= 109{,}797 \text{ MB} \equiv 110 \text{ GB}$$

A savings of 21.2% in staging disk requirement can be obtained by adding 10 more 300-MFLOP processors and operating ideally at duty cycle 1. For more realistic duty cycles (less than 1), the staging disk requirement can be calculated similarly.

The CERES daily average theoretical I/O bandwidth at CPU is shown in Table 6.3-2a,b for Releases A and B, respectively. Subsystems 4 and 5 are key contributors to both theoretical Processor <--> Data Handler throughput and I/O bandwidth at CPU. The CERES daily average I/O bandwidth at CPU for host attached storage is 1.25 MB/s and 4.8 MB/s for Releases A and B, respectively. Similarly, the theoretical Processor <--> Data Handler throughput is 0.4251 MB/s and 1.61 MB/s for Releases A and B, respectively. The needed I/O bandwidths at CPU for Network Attached Storage is also listed. Note that these numbers are based on raw AHWGP data for Standard Processing only.

Table 6.3-2a. CERES Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release A LaRC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
CERES 1aT	31,451	257,961	307,647	0.0093	0.0284	0.0099
CERES 2aT	113,237	119,628	193,473	0.0075	0.0211	0.0062
CERES 2bT	40,983	40,983	42,213	0.0001	0.0002	0.0000
CERES 3aT	40,983	242,471	283,815	0.0003	0.0008	0.0003
CERES 4aF	107,315	88,503	164,600	0.1557	0.4422	0.1308
CERES 5aF	55,608	79,620	135,589	0.1075	0.3228	0.1078
CERES 5aV	55,608	79,620	135,589	0.0173	0.0521	0.0174
CERES 6aT	79,620	1,517	81,498	0.0645	0.1938	0.0648
CERES 6cT	102,441	3,033	3,033	0.0001	0.0002	0.0000
CERES 7aT	764,640	475,192	1,240,193	0.0079	0.0238	0.0080
CERES 8aT	2,851,151	132,303	2,983,815	0.0029	0.0086	0.0029
CERES 9aTF	55,608	758	56,727	0.0448	0.1347	0.0451
CERES 9bTF	758	758	1,517	0.0000	0.0000	0.0000
CERES 10aT	3,725,273	203,798	3,929,432	0.0038	0.0113	0.0038
CERES 11aT	32,895	32,895	65,790	0.0002	0.0006	0.0002
CERES 12aF	11,681	90,994	102,856	0.0033	0.0099	0.0033
Total	**22,521**	**14,207**	**34,592**	**0.4251**	**1.256**	**0.4004**

Table 6.3-2b. CERES Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC (1 of 2)

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
CERES 1aA	87	714	852	0.0093	0.0284	0.0099
CERES 1aT	87	714	852	0.0093	0.0284	0.0099
CERES 1bA	87	714	852	0.0093	0.0284	0.0099
CERES 2aA	314	331	536	0.0075	0.0211	0.0062
CERES 2aT	314	331	536	0.0075	0.0211	0.0062
CERES 2bA	3	3	4	0.0001	0.0002	0.0000
CERES 2bT	3	3	4	0.0001	0.0002	0.0000
CERES 3aA	3	20	24	0.0003	0.0008	0.0003
CERES 3aT	3	20	24	0.0003	0.0008	0.0003
CERES 3bTA	7	20	27	0.0003	0.0009	0.0003
CERES 4aF	7,371	6,078	11,305	0.1557	0.4422	0.1308
CERES 4bAF	67,798	9,898	77,721	0.8993	2.6981	0.8995
CERES 5aF	3,819	5,468	9,312	0.1075	0.3228	0.1078
CERES 5aV	616	882	1,502	0.0173	0.0521	0.0174
CERES 5cAF	3,819	5,468	9,312	0.1075	0.3228	0.1078
CERES 5cAV	616	882	1,502	0.0173	0.0521	0.0174
CERES 6aA	5,468	104	5,597	0.0645	0.1938	0.0648
CERES 6aT	5,468	104	5,597	0.0645	0.1938	0.0648
CERES 6cA	9	1	0	0.0001	0.0002	0.0000
CERES 6cT	9	0	0	0.0001	0.0002	0.0000
CERES 7aA	424	263	687	0.0079	0.0238	0.0080

CERES 7aT	424	263	687	0.0079	0.0238	0.0080
CERES 7c	85	39	124	0.0014	0.0043	0.0014

Table 6.3-2b. CERES Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC (2 of 2)

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
CERES 8aA	237	11	248	0.0029	0.0086	0.0029
CERES 8aT	237	11	248	0.0029	0.0086	0.0029
CERES 8c	237	11	248	0.0029	0.0086	0.0029
CERES 9aAF	3,819	52	3,896	0.0448	0.1347	0.0451
CERES 9aTF	3,819	52	3,896	0.0448	0.1347	0.0451
CERES 9bAF	0	0	0	0.0000	0.0000	0.0000
CERES 9bTF	0	0	0	0.0000	0.0000	0.0000
CERES 10aA	310	17	326	0.0038	0.0113	0.0038
CERES 10aT	310	17	326	0.0038	0.0113	0.0038
CERES 10bTA	356	17	373	0.0043	0.0130	0.0043
CERES 11aA	9	9	18	0.0002	0.0006	0.0002
CERES 11aT	9	9	18	0.0002	0.0006	0.0002
CERES 12aF	32	252	285	0.0033	0.0099	0.0033
Total	**106,209**	**32,783**	**136,940**	**1.608**	**4.80**	**1.59**

6.3.1.1.4.2 MISR

The relationship between duty cycle and number of processors (for different peak MFLOPS ratings) is illustrated in Table 6.3-3 for MISR based on CPU requirements shown in Table 5.3-3. Ideally, at 300-MFLOPS peak rating, we need 46.1 fully dedicated processors (at duty cycle = 1) to satisfy MISR CPU requirements. To be more realistic, this is equivalent to 57.6 processors at 80% duty cycle. Similarly, 23.1 processors with 600 MFLOPS peak rating at duty cycle = 1 (ideal case) is equivalent to 28.9 processors at 80% duty cycle. With 900 MFLOPS peak processors, we need only 3.9 processors (the difference of 19.3 and 15.4) to allow a 20% slack in duty cycle.

Table 6.3-3. MISR Daily Average Processor Requirements as a Function of Duty Cycle for Release B

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	46.1	23.1	15.4
0.95	48.5	24.3	16.2
0.90	51.2	25.7	17.1
0.85	54.2	27.2	18.2
0.80	57.6	28.9	19.3

For a duty cycle of unity, to compute the daily average staging disk requirement we multiply the volume at completion for each MISR process by the number of activations per day (see Table 5.3-3) and sum over all processes. The daily average staging disk requirement determined by this method yields 556,305 MB for MISR. This staging disk capacity is needed when 46.1 300-MFLOPS (peak) processors are used at duty cycle unity. The theoretical estimate of staging disk requirement also indicates that 46.1 300-MFLOPS (peak) and 15.4 900-MFLOPS (peak) processors are equivalent. This is true only in the static sense because dynamically, processes on a 900-MFLOPS processor complete faster than on 300- or 600-MFLOPS processors, thereby requiring shorter residency times in the disk.

From Table 5.3-3 Part 2, based on the number of CPU requirements per day for MISR, we can calculate the time required for completion of all processes. If we use 58 300-MFLOP processors (at duty cycle 1) instead of 46, we can increase the processing capacity and reduce the time it takes for all MISR processes to complete. Substituting the appropriate numbers in equations (3) and (2), respectively:

$$\textit{Time of completion} = 2.99 \times 10^{8} / (58 \times 300.0 \times 0.25 \times 3600 \times 24)$$

$$= 0.795 \textit{ days}$$

$$\textit{Staging disk requirement} = 556{,}305 \textit{ MB per day} \times 0.795 \textit{ day}$$

$$= 442{,}262 \textit{ MB} \equiv 443 \textit{ GB}$$

A savings of 20.5% in staging disk requirement can be attained by adding 6 more 1000-MFLOP processors and operating ideally at duty cycle 1. For more realistic duty cycles, the staging disk requirement can be calculated similarly.

As shown in Table 6.3-4, the daily average I/O bandwidth at CPU (for host attached storage) for all MISR processes is 19.31 MB/s. The theoretical Processor <--> Data Handler throughput is 6.44 MB/s. Again, recall that these numbers are calculated based on raw AHWGP data for Standard Processing only.

Table 6.3-4. MISR daily average theoretical I/O bandwidth at CPU for Release B at LaRC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
MISP1A	45,924	62,119	108,043	1.2505	3.7515	1.2505
MISP1B1	55,329	56,974	112,303	1.2998	3.8994	1.2998
MISP1B2	110,549	63,293	173,841	2.0120	6.0361	2.0120
MISP2AS	86,459	6,986	93,603	1.0815	3.2464	1.0834
MISP2TC	64,521	3,995	68,515	0.7930	2.3790	0.7930
Total	**362,781**	**193,365**	**556,305**	**6.44**	**19.31**	**6.43**

6.3.1.1.4.3 MOPITT

Table 6.3-5 gives MOPITT daily average processor requirements as a function of processor duty cycle based on CPU requirements given in Table 5.3-4. For MOPITT, a 20% slack in duty cycle can be accommodated with a single 300-MFLOPS peak processor.

Table 6.3-5. MOPITT Daily Average Processor Requirements as a Function of Duty Cycle

Duty cycle	Number of 300-MFLOP processors
1	0.24
0.95	0.25
0.90	0.27
0.85	0.28
0.80	0.30

For a duty cycle of unity, to compute the daily average staging disk requirement, we multiply the volume at completion for each MOPITT process by the number of activations per day (in Table 5.3-4) and sum over all processes. The daily average staging disk requirement determined by this method yields 33 GB for MOPITT. This staging disk capacity is needed when 0.7 100-MFLOP (peak) processors are used at duty cycle unity. MOPITT staging disk volumes are not of major concern. However, we should determine ways where we can let other instruments share disk resources with MOPITT to reduce cost in buying a separate disk.

The MOPITT theoretical Processor <--> Data Handler throughput and I/O bandwidth at CPU shown in Table 6.3-6 are small compared to MISR and CERES.

Table 6.3-6. MOPITT Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
MOPL1	255	101	359	0.0041	0.0124	0.0042
MOPL1Qi-D	101	10	111	0.0013	0.0039	0.0013
MOPL2-E	31,311	185	31,756	0.3645	1.0966	0.3675
MOPL2Qi-D	175	10	185	0.0021	0.0064	0.0021
MOPL3	74	13	87	0.0010	0.0030	0.0010
MOPL3Qi-F	11	1	13	0.0001	0.0004	0.0001
Total	**31,928**	**320**	**32,511**	**0.3732**	**1.12**	**0.376**

6.3.1.1.4.4 ACRIM

Table 6.3-7 lists the ACRIM daily average theoretical I/O bandwidth at CPU and Processor <--> Data Handler throughput. One processor for ACRIM is more than adequate.

Table 6.3-7. ACRIM Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at LaRC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
ACRIM 1A		45	45	0.0005	0.0016	0.0005

6.3.1.1.4.5 LIS

Table 6.3-8 shows the LIS daily average theoretical I/O bandwidth at CPU and Processor <--> Data Handler throughput. One processor for LIS is more than adequate for both Releases A and B.

Table 6.3-8. LIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Releases A and B at MSFC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
LIS	86	1,256	1,325	0.0299	0.0464	0.0153

6.3.1.1.4.6 ASTER

Table 6.3-9 shows the daily average processor requirements as a function of duty cycle for ASTER. Two processors may be sufficient for ASTER processing. From Table 5.3-7, the number of MFPOs/day for ASTER is 8.57E6 for all processes. If three 300-MFLOP processors are used:

$$\begin{aligned} \text{Time of completion} &= 8.57 \times 10^6 / (3 \times 300.0 \times 0.25 \times 3600 \times 24) \\ &= 0.44 \text{ days} \end{aligned}$$

For estimating staging disk requirement, the total volume at completion of all processes is 1,154,487 MB.

$$\begin{aligned} \text{Staging disk requirement} &= 1{,}154{,}487 \text{ MB per day } \times 0.44 \text{ day} \\ &= 507{,}974 \text{ MB } \equiv 507 \text{ GB} \end{aligned}$$

Table 6.3-9. ASTER Daily Average Processor Requirements as a Function of Duty Cycle for Release B

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	1.32	0.66	0.44
0.95	1.39	0.69	0.46
0.90	1.46	0.73	0.49
0.85	1.55	0.77	0.52
0.80	1.65	0.83	0.55

The daily average theoretical I/O bandwidth at CPU and the Processor <--> Data handler throughput for ASTER is shown in Table 6.3-10. When Network Attached Disks are used, the needed bandwidth is much lower.

Table 6.3-10. ASTER Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at EDC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
AST_PGE_01	44,226	3,276	62,235	0.5498	1.8199	0.7203
AST_PGE_02	47,502	12,212	65,538	0.6911	2.1408	0.7585
AST_PGE_03	17,010	420	17,475	0.2017	0.6057	0.2023
AST_PGE_04	53,507	33,292	109,813	1.0046	3.2802	1.2710
AST_PGE_05	54,767	12,740	90,520	0.7813	2.6104	1.0477
AST_PGE_06	1,680	560	2,492	0.0259	0.0807	0.0288
AST_PGE_07	9,156	504	9,660	0.1118	0.3354	0.1118
AST_PGE_08	9,156	1,008	10,668	0.1176	0.3588	0.1235
AST_PGE_09	1,518	108	2,810	0.0188	0.0702	0.0325
AST_PGE_10	243	35	278	0.0032	0.0097	0.0032
Total	**238,766**	**64,155**	**371,489**	**3.506**	**11.31**	**4.30**

6.3.1.1.4.7 MODIS at EDC

The relationship between duty cycle and number of processors (for different peak MFLOPS ratings) is illustrated in Table 6.3-11 for MODIS at EDC based on CPU requirements shown in Table 5.3-8. Ideally, at 300-MFLOPS peak rating, we need 17.9 fully dedicated processors (at duty cycle = 1) to satisfy MODIS CPU requirements. To be more realistic, this is equivalent to 22.4 processors at 80% duty cycle. Similarly, 8.9 processors with 600 MFLOPS peak rating at duty cycle = 1 (ideal case) is equivalent to 11.1 processors at 80% duty cycle. With 900 MFLOPS peak processors, we need only 1.5 processors (the difference of 7.4 and 5.9) to allow a 20% slack in duty cycle.

Table 6.3-11. MODIS (EDC) Daily Average Processor Requirements as a Function of Duty Cycle for Release B

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	17.9	8.9	5.9
0.95	18.8	9.4	6.2
0.90	19.9	9.9	6.6
0.85	21.1	10.5	6.9
0.80	22.4	11.1	7.4

For a duty cycle of unity, to compute the daily average staging disk requirement we multiply the volume at completion for each MODIS process by the number of activations per day (see Table 5.3.8) and sum over all processes. The daily average staging disk requirement determined by this method yields 5,737,245 MB for MODIS (at EDC). This staging disk capacity is needed when 17.9 300- MFLOPS (peak) processors are used at duty cycle unity. The theoretical estimate of staging disk requirement also indicates that 17.9 300-MFLOPS (peak) and 5.9 900-MFLOPS (peak) processors are equivalent. This is true only in the static sense because dynamically, processes on a 900-MFLOPS processor complete faster than on 300- or 600-MFLOPS processors, thereby requiring shorter residency times in the disk.

From Table 5.3-8 Part 2, based on the number of CPU requirements per day for MODIS, we can calculate the time required for completion of all processes. If we use 23 300-MFLOP processors (at duty cycle 1) instead of 17.9, we can increase the processing capacity and reduce the time it takes for all MODIS processes to complete. Substituting the appropriate numbers in equations (3) and (2), respectively:

$$\textit{Time of completion} = 1.16 \times 10^{8} / (23 \times 300.0 \times 0.25 \times 3600 \times 24)$$

$$= 0.778 \textit{ days}$$

$$\textit{Staging disk requirement} = 5{,}737{,}245 \textit{ MB per day} \times 0.778 \textit{ day}$$

$$= 4{,}463{,}576 \textit{ MB} \equiv 4464 \textit{ GB}$$

The calculated theoretical requirement is an overestimate because the static model does not take into account disk space released by lower level processes that have completed before higher level processes begin execution. As shown in Table 6.3-12, the daily average I/O bandwidth at CPU for host attached storage for all EDC MODIS processes is 194.9 MB/s. Network Attached Storage appears promising. The theoretical Processor <--> Data Handler throughput is 65.14 MB/s. Again, recall that these numbers are calculated based on raw AHWGP data for Standard Processing only.

Table 6.3-12. MODIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at EDC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
MOD09:L3:9DY:G	2,572,271	42,600	2,585,564	30.2647	90.4549	29.9255
MOD12:L3:3MN:I	1,920,900	112	1,900,098	22.2339	66.4598	21.9919
MOD14:L3:10DY:G	959	96	1,054	0.0122	0.0366	0.0122
MOD14:L3:DY:G	10,544	959	11,502	0.1331	0.3994	0.1331
MOD14:L3:MN:G	96	32	128	0.0015	0.0044	0.0015
MOD15:L4:WK:G	212,310	152	213,274	2.4591	7.3866	2.4684
MOD16:L3:WK:G	761	4,382	5,143	0.0595	0.1786	0.0595
MOD17:L4:WK:G	2,505	8,763	11,269	0.1304	0.3913	0.1304
MOD34:L3:10DY:I	287,550	511	288,416	3.3340	10.0062	3.3382
MOD34:L3:MN:I	511	1,022	1,534	0.0178	0.0533	0.0178
MOD40:L3:10DY:I	305	31	336	0.0039	0.0117	0.0039
MOD40:L3:DY:I	560,723	305	561,028	6.4934	19.4801	6.4934
MOD40:L3:MN:I	31	10	41	0.0005	0.0014	0.0005
Total	**5,569,465**	**58,975**	**5,579,386**	**65.14**	**194.9**	**64.57**

6.3.1.1.4.8 COLOR

One processor is more than adequate for COLOR as shown in Table 6.3-13.

Table 6.3-13. COLOR Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	0.47	0.23	0.15
0.95	0.49	0.24	0.16
0.90	0.52	0.25	0.17

0.85	0.55	0.27	0.18
0.80	0.59	0.29	0.19

Table 6.3-14 shows the COLOR daily average theoretical I/O bandwidth at CPU and Processor <--> Data Handler throughput.

Table 6.3-14. COLOR Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
CLR Aero Browse	720	1	721	0.0083	0.0250	0.0083
CLR Chl. Browse	720	1	721	0.0083	0.0250	0.0083
CLR K 490 Browse	720	1	721	0.0083	0.0250	0.0083
CLR Rad. Browse	720	9	729	0.0084	0.0253	0.0084
CLR WLR Browse	720	1	721	0.0083	0.0250	0.0083
COLOR 1B	612	699	1,310	0.0152	0.0455	0.0152
COLOR Aerosol	720	37	757	0.0088	0.0263	0.0088
COLOR Ancillary	720	1	721	0.0083	0.0250	0.0083
COLOR Binned	720	387	1,107	0.0128	0.0384	0.0128
COLOR Chloro-A	720	37	757	0.0088	0.0263	0.0088
COLOR GAC	699	466	1,165	0.0135	0.0404	0.0135
COLOR K 490	720	37	757	0.0088	0.0263	0.0088
COLOR W.L. Rad.	720	37	757	0.0088	0.0263	0.0088
Total	**9230**	**1714**	**10,944**	**0.1267**	**0.38**	**0.1267**

6.3.1.1.4.9 MODIS at GSFC

The relationship between duty cycle and number of processors (for different peak MFLOPS ratings) is illustrated in Table 6.3-15 for MODIS at GSFC based on CPU requirements shown in Table 5.3-6. Ideally, at 300-MFLOPS peak rating, we need 29.6 fully dedicated processors (at duty cycle = 1) to satisfy MODIS CPU requirements. To be more realistic, this is equivalent to

37 processors at 80% duty cycle. Similarly, 14.8 processors with 600 MFLOPS peak rating at duty cycle = 1 (ideal case) is equivalent to 18.5 processors at 80% duty cycle. With 900 MFLOPS peak processors, we need only 2.5 processors (the difference of 12.4 and 9.9) to allow a 20% slack in duty cycle.

Table 6.3-15. MODIS Daily Average Processor Requirements as a Function of Duty Cycle for Release B at GSFC

Duty cycle	Number of 300-MFLOP processors	Number of 600-MFLOP processors	Number of 900-MFLOP processors
1	29.6	14.8	9.9
0.95	31.1	15.6	10.4
0.90	32.9	16.4	11.0
0.85	34.8	17.4	11.6
0.80	37.0	18.5	12.4

For a duty cycle of unity, to compute the daily average staging disk requirement we multiply the volume at completion for each MODIS process by the number of activations per day (see Table 5.3-6) and sum over all processes. The daily average staging disk requirement determined by this method yields 4,436,064 MB for MODIS (at GSFC). This staging disk capacity is needed when 29.6 300- MFLOPS (peak) processors are used at duty cycle unity. The theoretical estimate of staging disk requirement also indicates that 29.6 300-MFLOPS (peak) and 9.9 900-MFLOPS (peak) processors are equivalent. This is true only in the static sense because dynamically, processes on a 900-MFLOPS processor complete faster than on 300- or 600-MFLOPS processors, thereby requiring a shorter residency times in the disk.

From Table 5.3-6 Part 4, based on the number of CPU requirements per day for MODIS, we can calculate the time required for completion of all processes. If we use 40 300-MFLOP processors (at duty cycle 1) instead of 29.6, we can increase the processing capacity and reduce the time it takes for all MODIS processes to complete. Substituting the appropriate numbers in equations (3) and (2), respectively:

$$\textit{Time of completion} = 1.92 \times 10^{8} / (40 \times 300.0 \times 0.25 \times 3600 \times 24)$$

$$= 0.740 \textit{ days}$$

$$\textit{Staging disk requirement} = 4{,}436{,}064 \textit{ MB per day} \times 0.740 \textit{ day}$$

$$= 3{,}282{,}687 \textit{ MB} \equiv 3283 \textit{ GB}$$

The calculated theoretical requirement is an overestimate because the static model does not take into account disk space released by lower level processes that have completed before higher level processes begin execution. As shown in Table 6.3-16, the daily average I/O bandwidth at CPU for host attached storage for all GSFC MODIS processes is 112 MB/s. Network Attached Storage appears promising. The theoretical Processor <--> Data Handler throughput is 37.23 MB/s. Again, recall that these numbers are calculated based on raw AHWGP data for Standard Processing only.

Table 6.3-16. MODIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at GSFC (1 of 2)

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
MOD01:L1A:G	70,200	115,099	185,299	2.1447	6.4340	2.1447
MOD02:L1B:G	115,099	180,145	295,244	3.4172	10.2515	3.4172
MOD03:L1A:G	115,099	13,455	26,651	1.4879	3.2842	0.3085
MOD04:L2:G	109,582	1,091	111,258	1.2809	3.8496	1.2877
MOD04:L3:DY:G	2,081	175	2,256	0.0261	0.0783	0.0261
MOD04:L3:MN:G	164	41	205	0.0024	0.0071	0.0024
MOD04:L3:WK:G	1,324	175	1,500	0.0174	0.0521	0.0174
MOD05:L2:G	98,385	11,089	111,229	1.2671	3.8215	1.2874
MOD06:L2:G	291,383	10,249	325,792	3.4911	10.7530	3.7707
MOD06:L3:MN:G	63,855	0	63,855	0.7391	2.2172	0.7391
MOD09:L2:I	100,755	41,243	144,922	1.6435	4.9643	1.6773
MOD10:L2:I	61,021	410	64,356	0.7110	2.1669	0.7449
MOD11:L2:I	236,662	6,318	255,908	2.8123	8.5864	2.9619
MOD11:L3:WK:G	41,626	609	42,235	0.4888	1.4665	0.4888
MOD13:L2:G	133,991	15,795	152,711	1.7336	5.2348	1.7675
MOD14:L2:G	342,839	1,580	344,419	3.9863	11.9590	3.9863
MOD28:D:ORBIT:G	406	820	1,226	0.0142	0.0426	0.0142
MOD28:L2:G	197,572	3,159	200,731	2.3233	6.9698	2.3233
MOD28:L3:COMP:D:DY:G	875	396	1,271	0.0147	0.0441	0.0147
MOD28:L3:COMP:N:DY:G	875	396	1,271	0.0147	0.0441	0.0147
MOD28:L3:D:WK:G	444	444	888	0.0103	0.0308	0.0103
MOD28:L3:N:WK:G	444	444	888	0.0103	0.0308	0.0103

MOD28:L3:TMP:D:WK:G	499	111	610	0.0071	0.0212	0.0071
MOD28:L3:TMP:N:WK:G	1,019	227	1,246	0.0144	0.0433	0.0144

Table 6.3-16. MODIS Daily Average Theoretical I/O Bandwidth at CPU and Processor <--> Data Handler Throughput for Release B at GSFC (2 of 2)

Process	Volume staged (MB/day)	Volume destaged (MB/day)	I/O (MB/day)	Theoretical Processor <--> Data Handler Throughput (MB/s)	Needed I/O BW at CPU for host attached disk (MB/s)	Needed I/O BW at CPU for Network Attached Storage (MB/s)
MOD28:N:ORBIT:G	406	820	1,226	0.0142	0.0426	0.0142
MOD28:SPBIN:G	3,159	792	3,951	0.0457	0.1372	0.0457
MOD29:L2:G	47,566	790	48,356	0.5597	1.6790	0.5597
MOD32:L2:G	487	6	498	0.0057	0.0172	0.0058
MOD35:L2:G	208,810	3,171	289,493	2.4535	8.2576	3.3506
MOD41:L2:H	13,929	23,693	37,914	0.4354	1.3097	0.4388
MOD:ATMOS:L2:G	202,960	21,809	237,931	2.6015	7.9568	2.7538
MOD:ATMOS:L3:MN:G	5,431	0	5,431	0.0629	0.1886	0.0629
MODOCCLR:L2:G	96,332	49,108	145,586	1.6833	5.0517	1.6850
MODOCCLR:L3:COMP:DY:G	13,994	12,682	26,676	0.3087	0.9262	0.3087
MODOCCLR:L3:TMP:WK:G	14,681	3,623	7,247	0.2119	0.5076	0.0839
MODOCCLR:L3:WK:QC:G	14,494	14,494	28,987	0.3355	1.0065	0.3355
MODOCCLR:ORBIT:G	6,065	12,244	18,309	0.2119	0.6357	0.2119
MODOCCLR:SPBIN:G	49,108	6,336	55,443	0.6417	1.9251	0.6417
Total	**2,663,621**	**553,037**	**3,243,020**	**37.23**	**112**	**37.5**

6.3.1.2 Dynamic Analysis Using ECS Systems Performance Model

The ECS Systems Performance Model is used to dynamically simulate instrument processing at each of the DAACs for the "one instrument's products on one cluster" optimization alternative. The AHWGP data (January 1995 baseline) served as input to the model.

6.3.1.2.1 Assumptions

The following assumptions are made for the dynamic model simulations:

- Networks are not constrained;

- Disk storage is not constrained;
- No reprocessing;
- Data sets are not organized in the archive;
- No waiting storage (waiting storage may lower network needs);
- V0 migration is considered;
- User pull load is included based on user characterization data

6.3.1.2.2 CERES

As indicated in Table 5.3-2, CERES processing can be considered as episodic. Processes are activated 24 times a day, some once a day, once a week and once a month. The ECS Systems Performance Model is used to simulate CERES processing to make a dynamic assessment of process completion times, processing resource usage and Processing <--> Data Handler throughput. Two processor speeds (300- and 900-MFLOPS) are considered for the analyses. A processor efficiency factor of 0.25 is used for the simulation. The model simulation is made for a 3-week period. This 3-week period includes daily, weekly and monthly process activations on certain days.

6.3.1.2.2.1 CERES process completion times

There are 30 CERES processes from various subsystems. Table 6.3-17a illustrates the model simulated process completion times (maximum, minimum, average and the standard deviation) for each CERES process based on 300-MFLOPS processors. On the average, processes belonging to CERES Subsystems 4 and 5 take 300-3800 minutes to complete. Similarly, Table 6.3-17b illustrates model simulated process completion times for 900 MFLOPS.

Table 6.3-17a. CERES Process Completion Times for 300-MFLOPS

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
CERES 1aA	4.62	4.62	4.62	0.00
CERES 1aT	4.62	4.62	4.62	0.00
CERES 1bA	4.62	4.62	4.62	0.00
CERES 2aA	0.84	0.84	0.84	0.00
CERES 2aT	0.84	0.84	0.84	0.00
CERES 2bA	0.13	0.13	0.13	0.00
CERES 2bT	0.13	0.13	0.13	0.00
CERES 3aA	10.50	10.50	10.50	0.00
CERES 3aT	10.50	10.50	10.50	0.00
CERES 3bTA	21.00	21.00	21.00	0.00
CERES 4aF	7.56	7.56	7.56	0.00
CERES 4bAF	1271.18	756.00	765.03	67.50
CERES 5aF	593.88	593.88	593.88	0.00
CERES 5aV	593.88	593.88	593.88	0.00
CERES 5cAF	593.88	593.88	593.88	0.00
CERES 5cAV	593.88	593.88	593.88	0.00
CERES 6aA	3.75	1.09	1.11	0.17
CERES 6aT	1.09	1.09	1.09	0.00
CERES 6cA	6.52	6.52	6.52	0.00
CERES 6cT	6.52	6.52	6.52	0.00
CERES 7aA	151.20	151.20	151.20	0.00
CERES 7aT	151.20	151.20	151.20	0.00
CERES 7c	302.40	302.40	302.40	0.00
CERES 8aA	50.40	50.40	50.40	0.00
CERES 8aT	50.40	50.40	50.40	0.00
CERES 8c	100.80	100.80	100.80	0.00
CERES 9aAF	1.09	1.09	1.09	0.00
CERES 9aTF	1.09	1.09	1.09	0.00
CERES 9bAF	3.26	3.26	3.26	0.00
CERES 9bTF	3.26	3.26	3.26	0.00
CERES 10aA	54.60	54.60	54.60	0.00
CERES 10aT	54.60	54.60	54.60	0.00
CERES 10bTA	109.20	109.20	109.20	0.00
CERES 11aA	24.85	24.85	24.85	0.00
CERES 11aT	24.85	24.85	24.85	0.00
CERES 12aF	8.40	8.40	8.40	0.00

Table 6.3-17b CERES Process Completion Times for 900-MFLOPS (1 of 2)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
CERES 1aA	1.54	1.54	1.54	0.00
CERES 1aT	1.54	1.54	1.54	0.00
CERES 1bA	1.54	1.54	1.54	0.00
CERES 2aA	0.60	0.60	0.60	0.00
CERES 2aT	0.60	0.60	0.60	0.00
CERES 2bA	0.13	0.13	0.13	0.00
CERES 2bT	0.13	0.13	0.13	0.00
CERES 3aA	3.50	3.50	3.50	0.00
CERES 3aT	3.50	3.50	3.50	0.00
CERES 3bTA	7.00	7.00	7.00	0.00
CERES 4aF	197.41	2.52	3.17	9.11
CERES 4bAF	827.01	382.93	530.58	107.95
CERES 5aF	395.89	197.96	270.21	93.71
CERES 5aV	395.45	199.18	310.60	79.51
CERES 5cAF	454.23	197.96	268.08	93.22
CERES 5cAV	423.32	197.96	302.60	79.90
CERES 6aA	451.35	57.41	237.89	118.58
CERES 6aT	192.94	0.36	13.80	38.94
CERES 6cA	6.52	6.52	6.52	0.00
CERES 6cT	121.80	121.80	121.80	0.00
CERES 7aA	50.40	50.40	50.40	0.00
CERES 7aT	50.40	50.40	50.40	0.00
CERES 7c	100.80	100.80	100.80	0.00
CERES 8aA	16.80	16.80	16.80	0.00
CERES 8aT	16.80	16.80	16.80	0.00

Table 6.3-17b CERES Process Completion Times for 900-MFLOPS (2 of 2)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
CERES 8c	33.60	33.60	33.60	0.00
CERES 9aAF	256.87	0.36	76.95	94.98
CERES 9aTF	248.68	0.36	117.45	77.71
CERES 9bAF	115.05	115.05	115.05	0.00
CERES 9bTF	149.96	149.96	149.96	0.00
CERES 10aA	18.20	18.20	18.20	0.00
CERES 10aT	149.53	149.53	149.53	0.00
CERES 10bTA	54.13	54.13	54.13	0.00
CERES 11aA	406.27	227.63	316.95	89.32
CERES 11aT	109.18	24.85	67.01	42.16
CERES 12aF	139.99	2.80	10.42	31.42

6.3.1.2.2.2 CERES Processing Resource Usage

The episodic nature of CERES processing is clearly evident in the trace of Processing resource usage shown in Figure 6.3-2a,b for 300-MFLOPS processors. For a maximum of 48 processors (as shown in Table 5.3-2), CERES processing resources (CPU and disk) show clear peaks when daily, and/or weekly and/or monthly processes coincide. The periods of decreased activity can be varied by constraining in the model the maximum number of processors available. Reducing the maximum number of available processors will increase CPU loading and decrease the periods of reduced activity. These periods of reduced activity is known as "processing slack". A slack is essential during real time operations to account for sudden and unexpected down times. Similarly, Figure 6.3-3a,b illustrates processing resource usage for 900-MFLOPS.

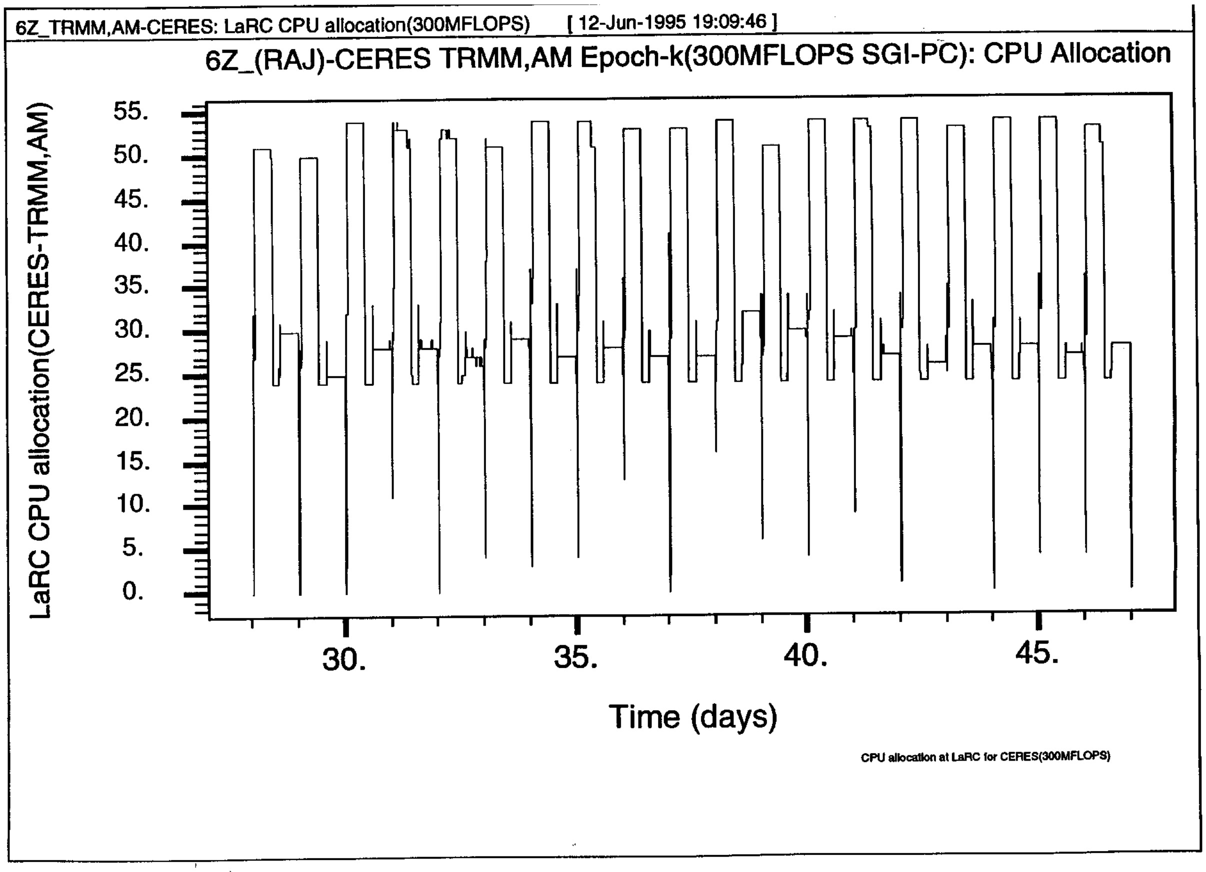

Figure 6.3-2a. CERES Processing Resource Usage - Number of CPUs (300 MFLOPS)

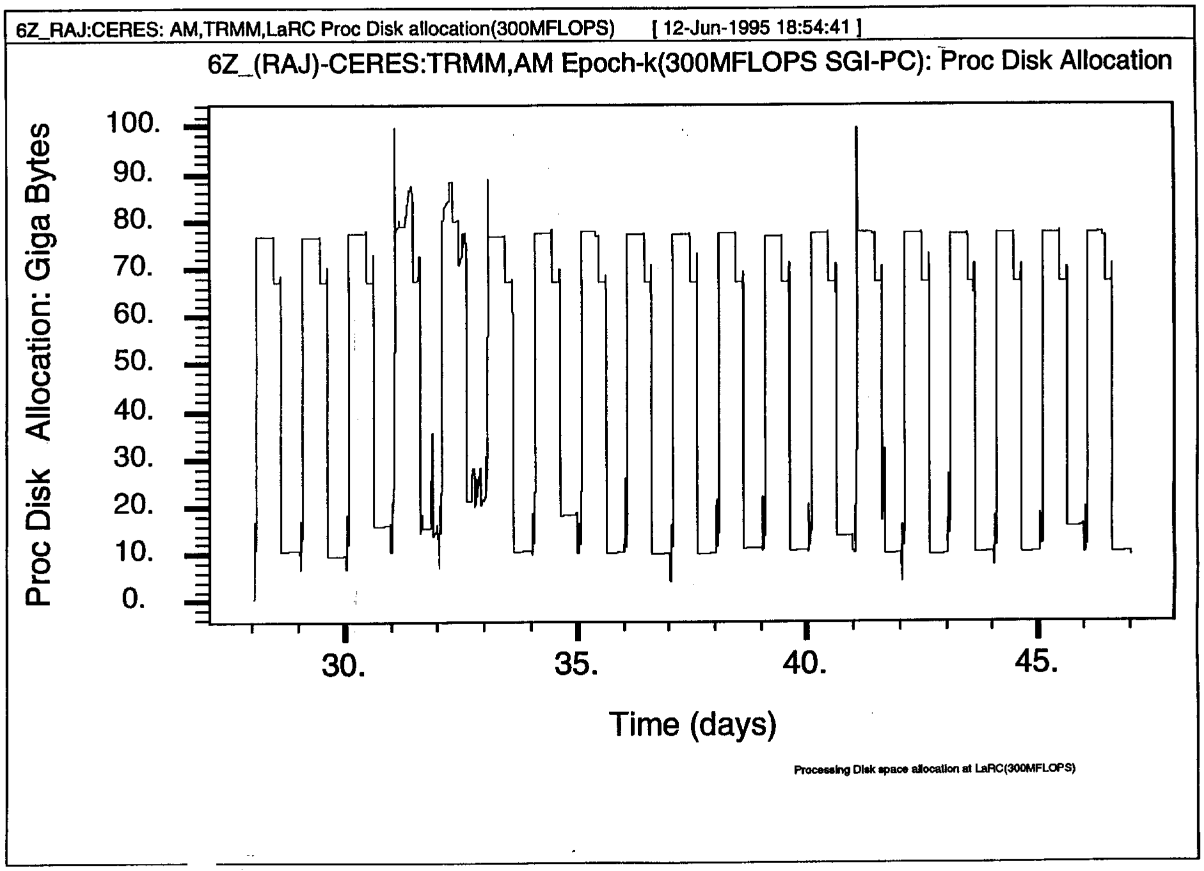

Figure 6.3-2b. CERES Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

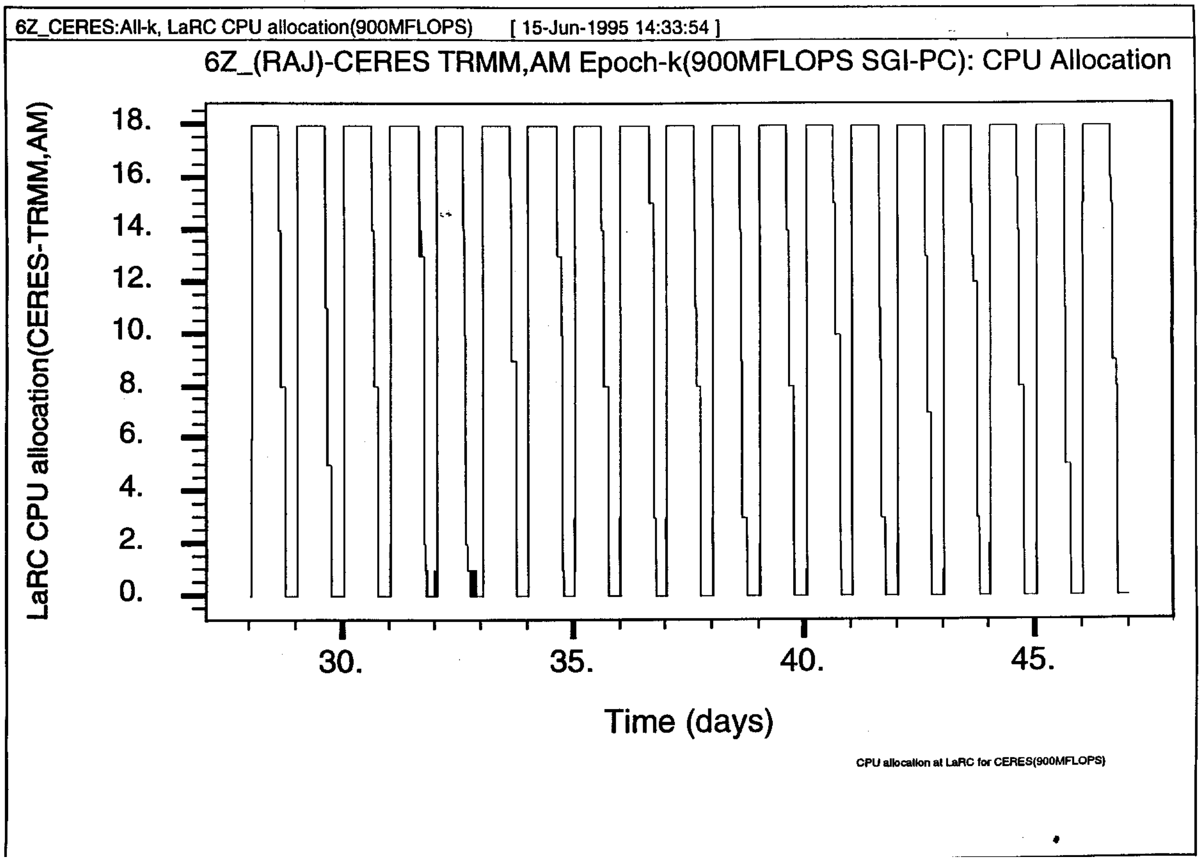

Figure 6.3-3a. CERES Processing Resource Usage - Number of CPUs (900 MFLOPS)

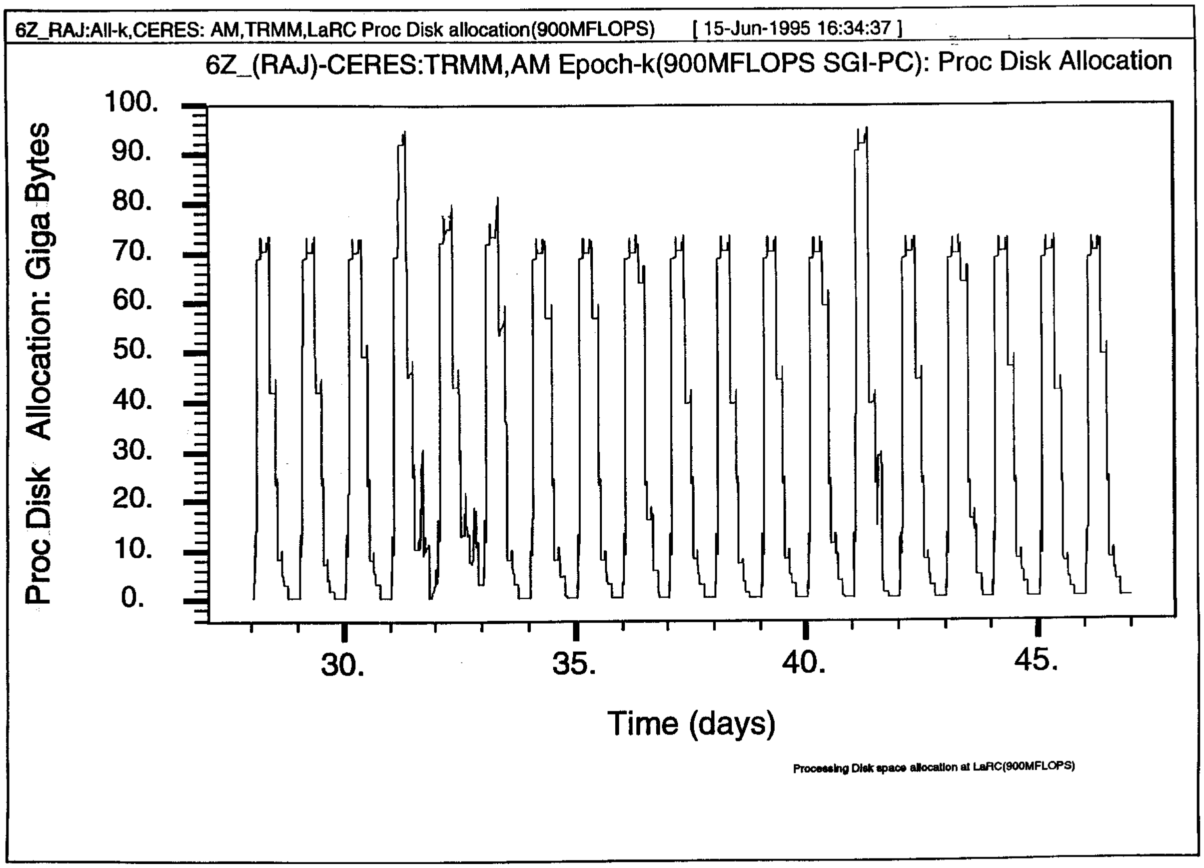

Figure 6.3-3b. CERES Processing Resource Usage - Processing Disk Capacity (900 MFLOPS)

Table 6.3-18a. CERES CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
36.0	54.0 (constrained)	46.0	100.1

Table 6.3-18b. CERES CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)

Number of CPUs		Staging disk capacity (GB)	
Average	Peak	Average	Peak
11.8	18.0 (constrained)	30.6	95.6

6.3.1.2.3 MISR

The MISR processing was dynamically simulated for a three week period. A 300-MFLOPS (peak) mid-range and 900-MFLOPS high-range processors with a processor efficiency factor of 0.25 were assumed for the simulation. Recall from Table 5.3-3 that MISR processing is performed by orbit (each process is activated approximately 14.56 times a day).

6.3.1.2.3.1 MISR Process Completion Times

Since MISR processes are activated on an average of 14.56 times a day, the processing is not episodic unlike CERES. Tables 6.3-19a,b show the process completion times for each MISR process. As shown in Table 6.3-19a (300 MFLOPS), the Level 1A and 1B1 processes (0 and 1) are estimated to run in less than 100 minutes. Process 2 (Level 1B2), however, can take on an average over 30 hours to complete. Processes 3 and 4 (Level 2 TC and AS) are estimated to take on an average 17 hours and 25 hours, respectively.

Table 6.3-19a. MISR Process Completion Times (300 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
MISP1A	52.67	52.66	52.67	0.00
MISP1B1	39.56	39.55	39.56	0.00
MISP1B2	1901.55	1901.55	1901.55	0.00
MISP2AS	1553.33	1553.33	1553.33	0.00
MISP2TC	1033.77	1033.77	1033.77	0.00

Table 6.3-19b. MISR Process Completion Times (900 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
MISP1A	90.79	17.55	27.16	19.62
MISP1B1	98.12	13.18	19.513	15.88
MISP1B2	683.58	633.85	635.09	5.31
MISP2AS	402.62	344.59	346.31	9.11
MISP2TC	564.00	517.77	521.43	9.38

6.3.1.2.3.2 MISR Processing Resource Usage

The simulated processing resource usage (number of CPUs and processing disk capacity) is illustrated in Figures 6.3-4a,b and 6.3-5a,b for 300 and 900 MFLOPS, respectively. Both input data staged for the L1 processes and the output produced by them exponentially fill up the processing disk. Higher level processes are activated only after all input data from lower level processes are available. Note that the number of CPUs required for processing gradually increases as higher level processes are activated. This simulation also represents a scenario where MISR processing is restarted after a total interruption. The simulation indicates that it can take up to three days for the production processing to attain steady state. Figure 6.3-4a and Table 6.3-20a indicate that an average of 46.4 processors (with 300 MFLOPS peak) are needed for MISR with an average staging disk volume of approximately 300 GB. With the processing load maintained by 46 processors, there are no jobs queued and waiting to be processed. MISR processing resource usage for 900 MFLOPS processors is shown in Figures 6.3-5a,b.

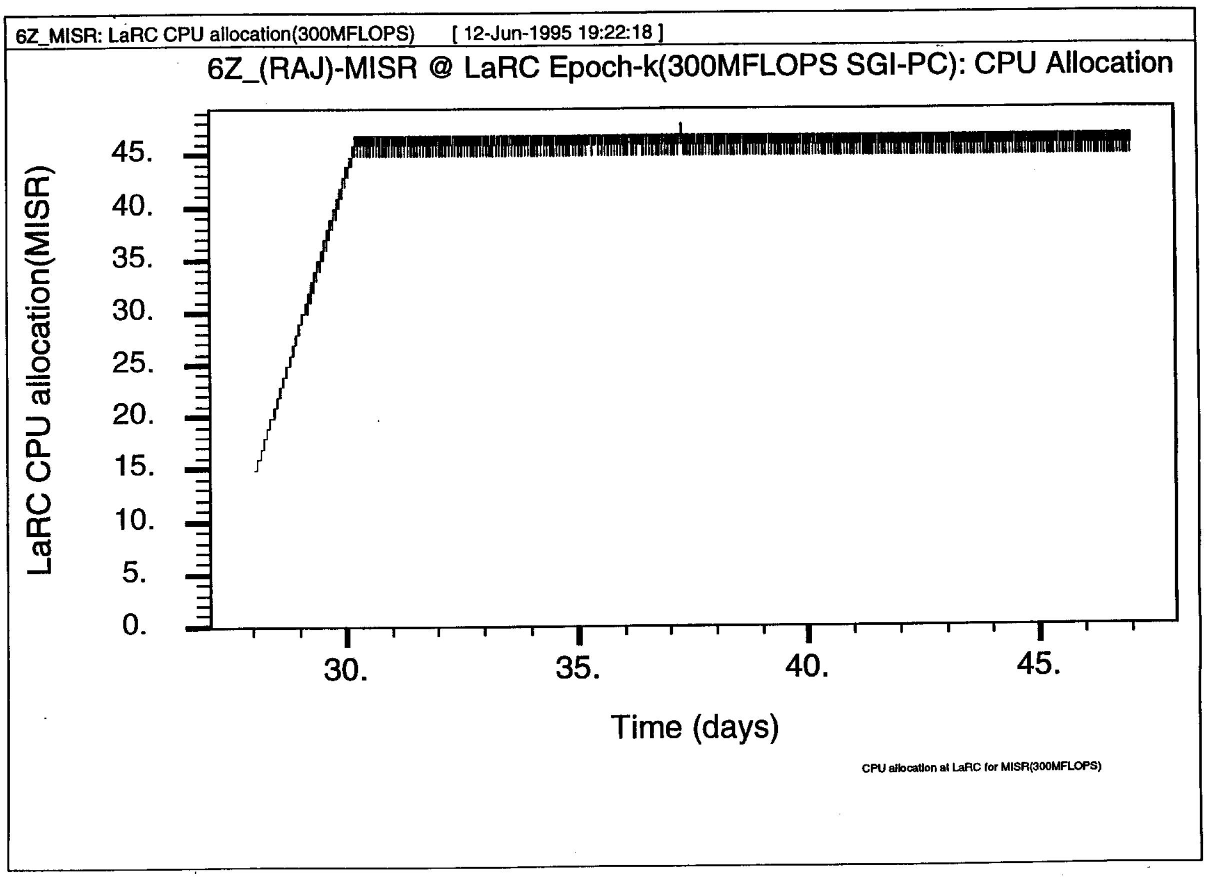

Figure 6.3-4a. MISR Processing Resource Usage - Number of CPUs (300 MFLOPS)

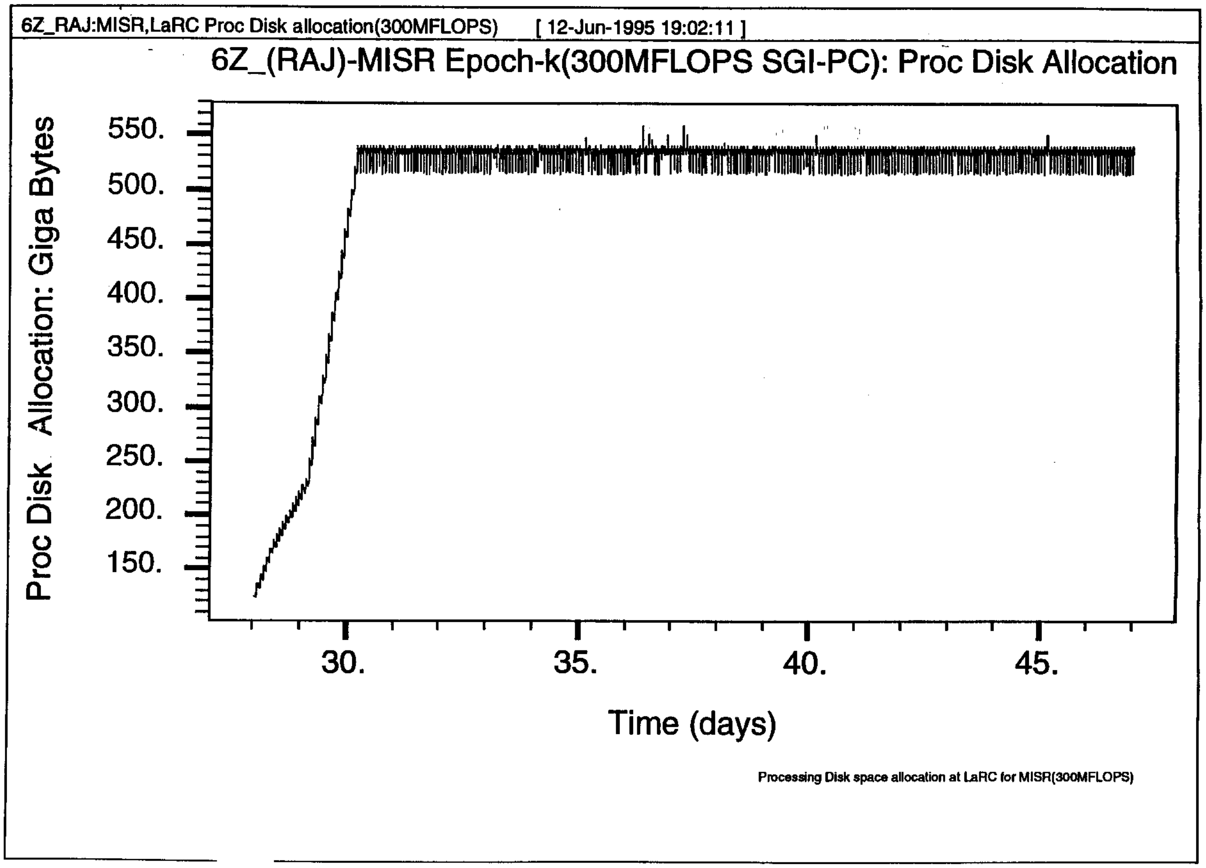

Figure 6.3-4b. MISR Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

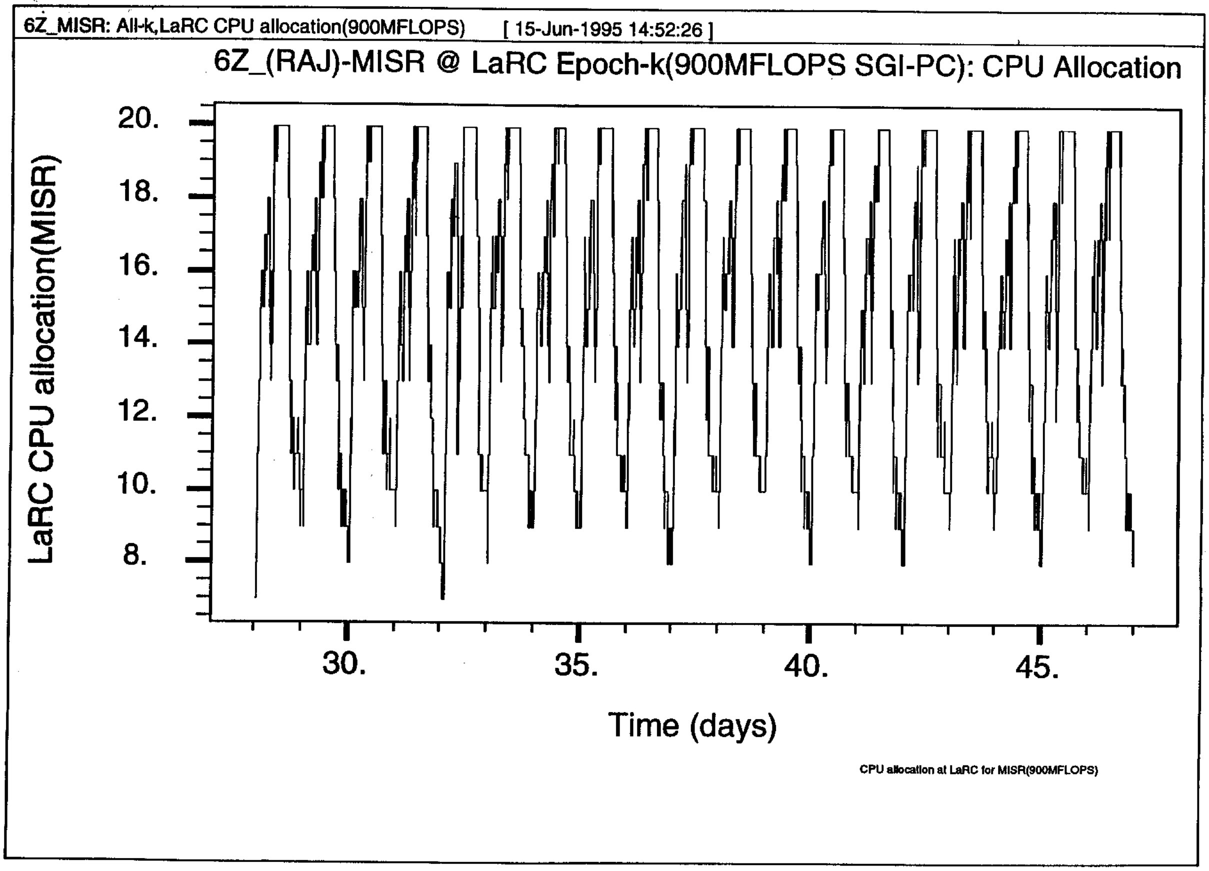

Figure 6.3-5a. MISR Processing Resource Usage - Number of CPUs (900 MFLOPS)

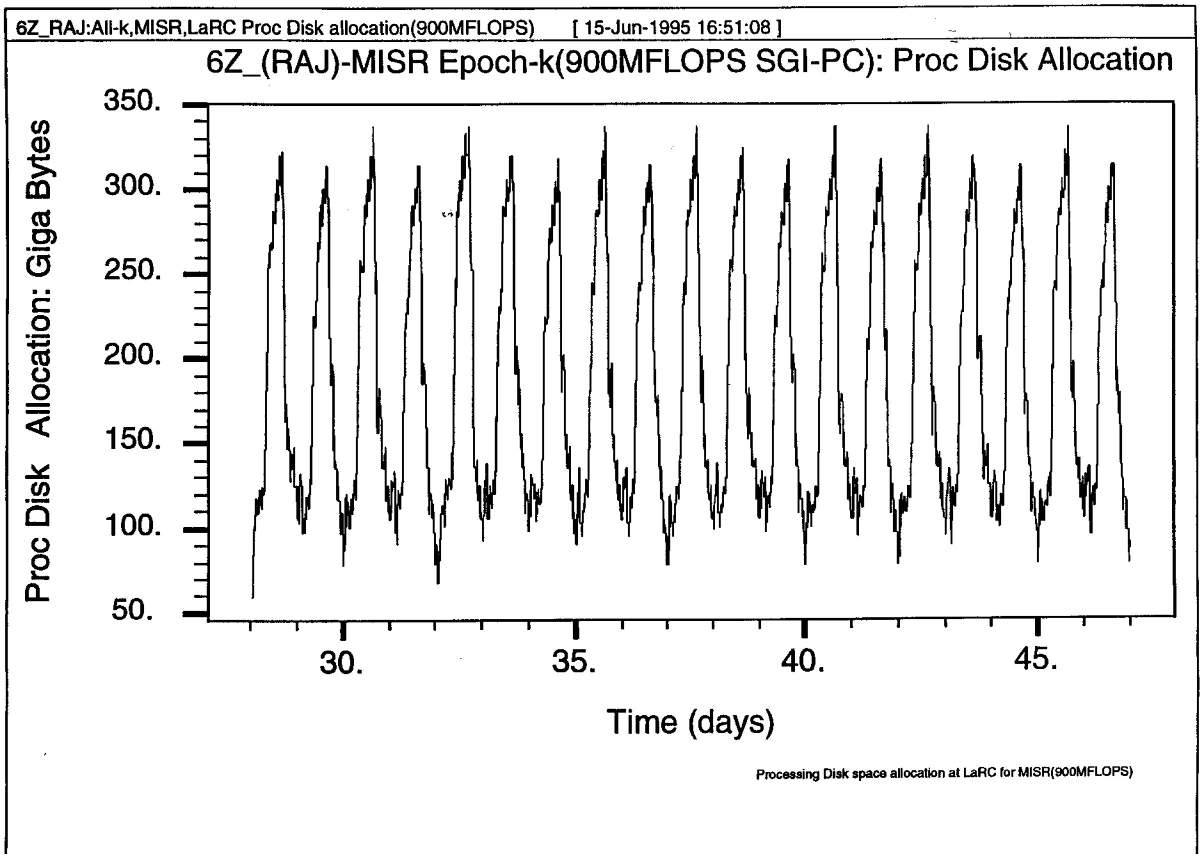

Figure 6.3-5b. MISR Processing Resource Usage - Processing Disk Capacity (900 MFLOPS)

Table 6.3-20a. MISR CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
46.4	48.0 (constrained)	534.87	560.34

Table 6.3-20b. MISR CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
15.4	20 (constrained)	186.4	338.3

6.3.1.2.4 MOPITT

The MOPITT L1 and L2 processes are activated once a day. The L3 processes are activated once a week. The simulation is performed for a 10-day period. This time period is representative of daily and weekly processing of all MOPITT products. A 100-MFLOP (peak) workstation-class processor with an efficiency factor of 0.25 is assumed for the simulation. The model run is constrained to 1 CPU.

6.3.1.2.4.1 MOPITT Process Completion Times

The process completion times of various MOPITT processes are illustrated in Table 6.3-21. With the exception of MOPL2-E (process number 2) which is estimated to take 17 hours, MOPITT processes are estimated to take less than 1 hour to complete.

Table 6.3-21. MOPITT Process Completion Times (300 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
MOPL1	3.73	3.73	3.73	0.00
MOPL1Qi-D	0.41	0.41	0.41	0.00
MOPL2-E	333.8	333.8	333.8	0.00
MOPL2Qi-D	0.30	0.30	0.30	0.00
MOPL3	6.29	6.29	6.29	0.00
MOPL3Qi-F	0.20	0.20	0.20	0.00

6.3.1.2.4.2 MOPITT Processing Resource Usage

The L1-L3 processing is episodic and coincides with the data arrival rates. Figures 6.3-6a,b illustrate MOPITT processing resource usage (number of CPUs and processing disk capacity). As shown in Table 6.3-22, the average daily processing disk space required is 9.7 GB with a peak requirement of 32.7 GB.

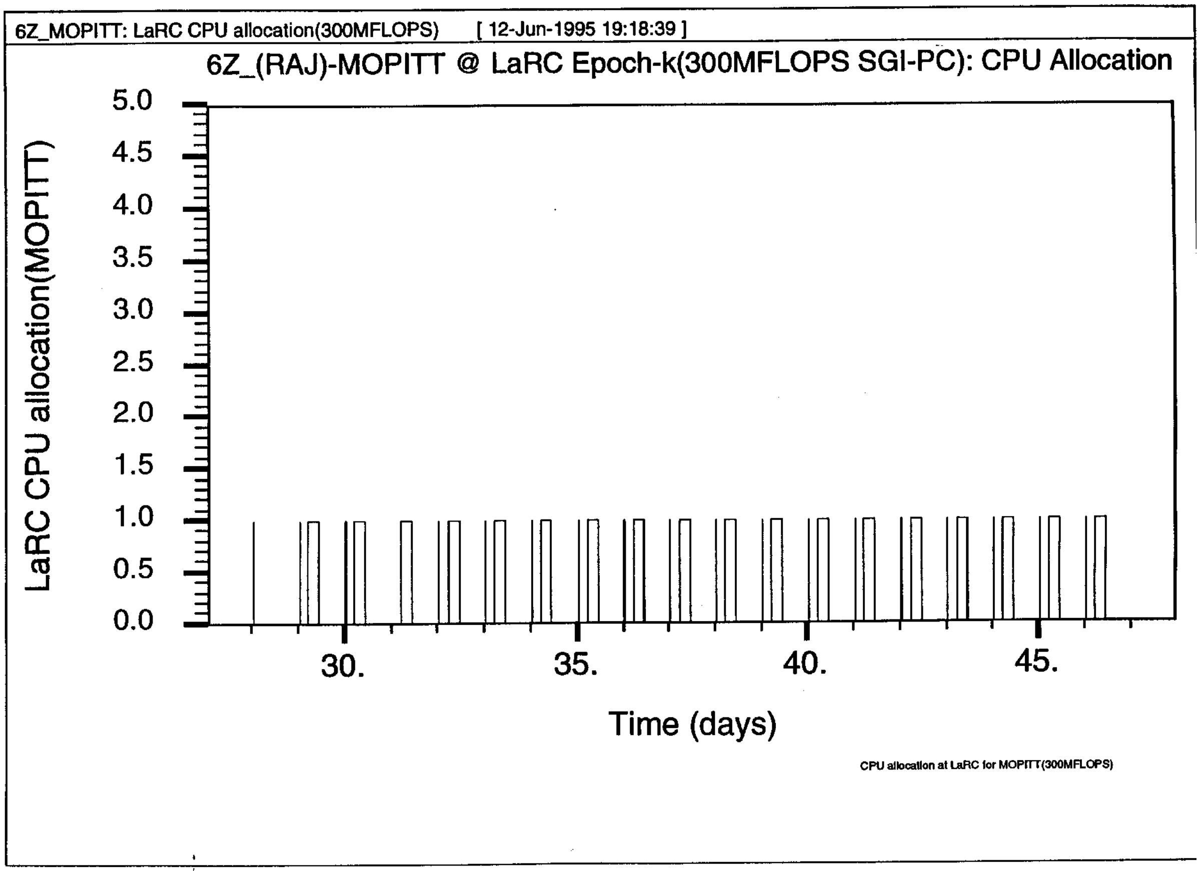

Figure 6.3-6a. MOPITT Processing Resource Usage - Number of CPUs (300 MFLOPS)

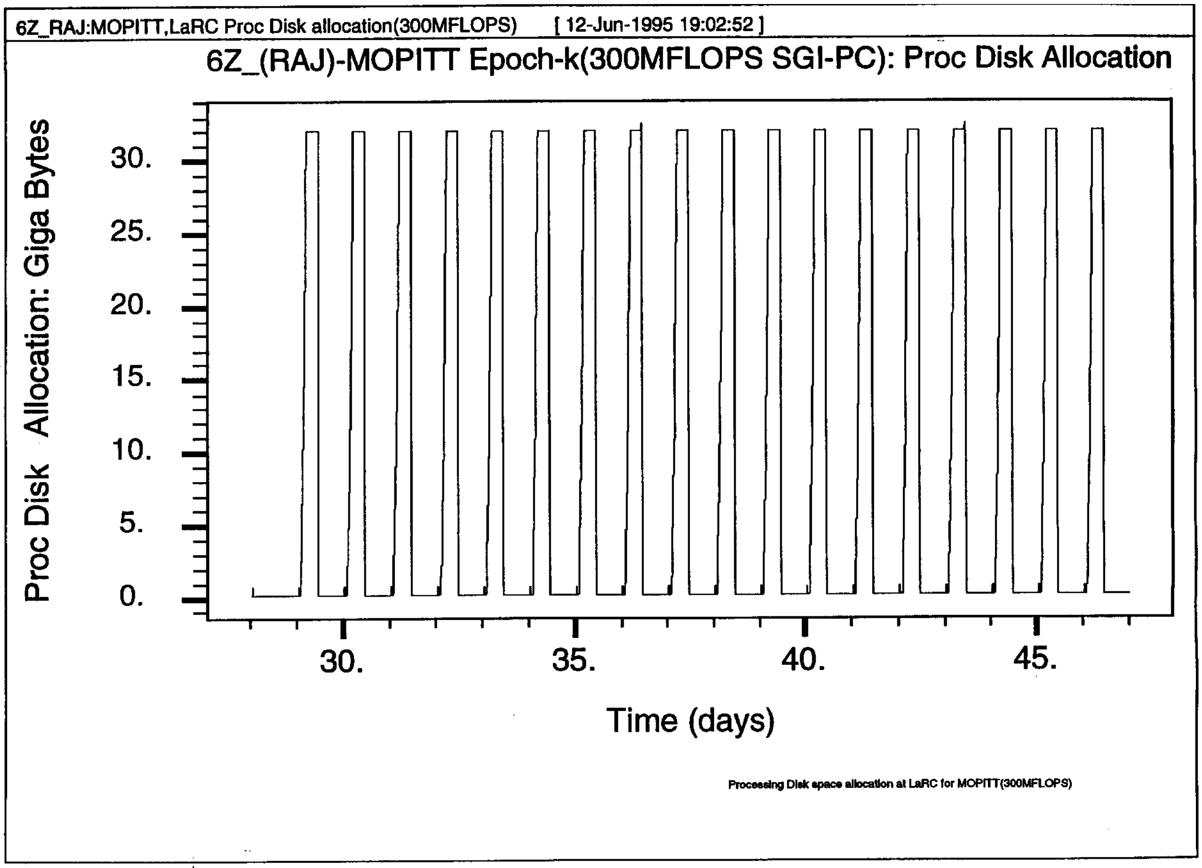

Figure 6.3-6b. MOPITT Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

Table 6.3-22. MOPITT CPU and Staging disk capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
0.23	1.0 (constrained)	9.7	32.7

6.3.1.2.5 MODIS at GSFC

The MODIS L3 processes have not been modeled dynamically due to uncertainties in the data. Therefore, the following results should be interpreted cautiously.

6.3.1.2.5.1 MODIS (GSFC) Process Completion Times

Tables 6.3-23a,b outline the process completion times for 300 and 900 MFLOPS, respectively. Refer to Table 5.3-6 for a summary of MODIS processing requirements.

Table 6.3-23a. MODIS (GSFC) Process Completion Times (300 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
MOD01:L1A:G	3.27	3.27	3.27	0.00
MOD02:L1B:G	42.63	42.63	42.63	0.00
MOD03:L1A:G	1.34	1.34	1.34	0.00
MOD04:L2:G	0.43	0.43	0.43	0.00
MOD05:L2:G	0.43	0.43	0.43	0.00
MOD06:L2G	2.37	2.37	2.37	0.00
MOD:ATMOS:L2:G	0.46	0.46	0.46	0.00
MOD35:L2:G	1.98	1.98	1.98	0.00
MOD11:L2:I	0.49	0.49	0.49	0.00
MOD09:L2:I	0.56	0.56	0.56	0.00
MOD10:L2:I	0.24	0.24	0.24	0.00
MOD14:L2:G	0.66	0.66	0.66	0.00
MOD13:L2:G	0.59	0.59	0.59	0.00
MOD29:L2:G	0.18	0.18	0.18	0.00
MOD41:L2:H	0.17	0.17	0.17	0.00
MODOCCLR:L2:G	18.37	18.37	18.37	0.00
MODOCCLR:SPBIN:G	3.67	3.67	3.67	0.00
MOD28:L2:G	6.77	6.77	6.77	0.00
MOD28:SPBIN:G	2.40	2.40	2.40	0.00

Table 6.3-23b. MODIS (GSFC) Process Completion Times (900 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
MOD01:L1A:G	3.10	1.09	1.10	0.09
MOD02:L1B:G	16.16	14.21	14.21	0.07
MOD03:L1A:G	2.48	0.44	0.45	0.07
MOD04:L2:G	2.64	0.43	0.47	0.23
MOD05:L2:G	2.48	0.43	0.45	0.17
MOD06:L2G	3.09	0.79	0.81	0.14
MOD:ATMOS:L2:G	2.50	0.46	0.47	0.09
MOD35:L2:G	2.70	0.66	0.67	0.10
MOD11:L2:I	1.73	0.49	0.50	0.05
MOD09:L2:I	2.51	0.56	0.57	0.08
MOD10:L2:I	1.15	0.24	0.25	0.04
MOD14:L2:G	1.58	0.66	0.67	0.04
MOD13:L2:G	2.96	0.59	0.61	0.17
MOD29:L2:G	1.12	0.18	0.19	0.04
MOD41:L2:H	1.27	0.14	0.16	0.07
MODOCCLR:L2:G	8.31	6.12	6.15	0.17
MODOCCLR:SPBIN:G	3.32	1.22	1.23	0.06
MOD28:L2:G	4.51	2.25	2.28	0.19
MOD28:SPBIN:G	3.12	0.80	0.83	0.20

6.3.1.2.5.2 MODIS (GSFC) Processing Resource Usage

Figures 6.3-7a,b and 6.3-8a,b illustrate dynamic simulation of number of CPUs, staging disk capacity for both 300 and 900 MFLOPS processors, respectively. The number of processors and staging disk capacities for 300 and 900 MFLOPS processors is listed in Tables 6.3.24-a,b.

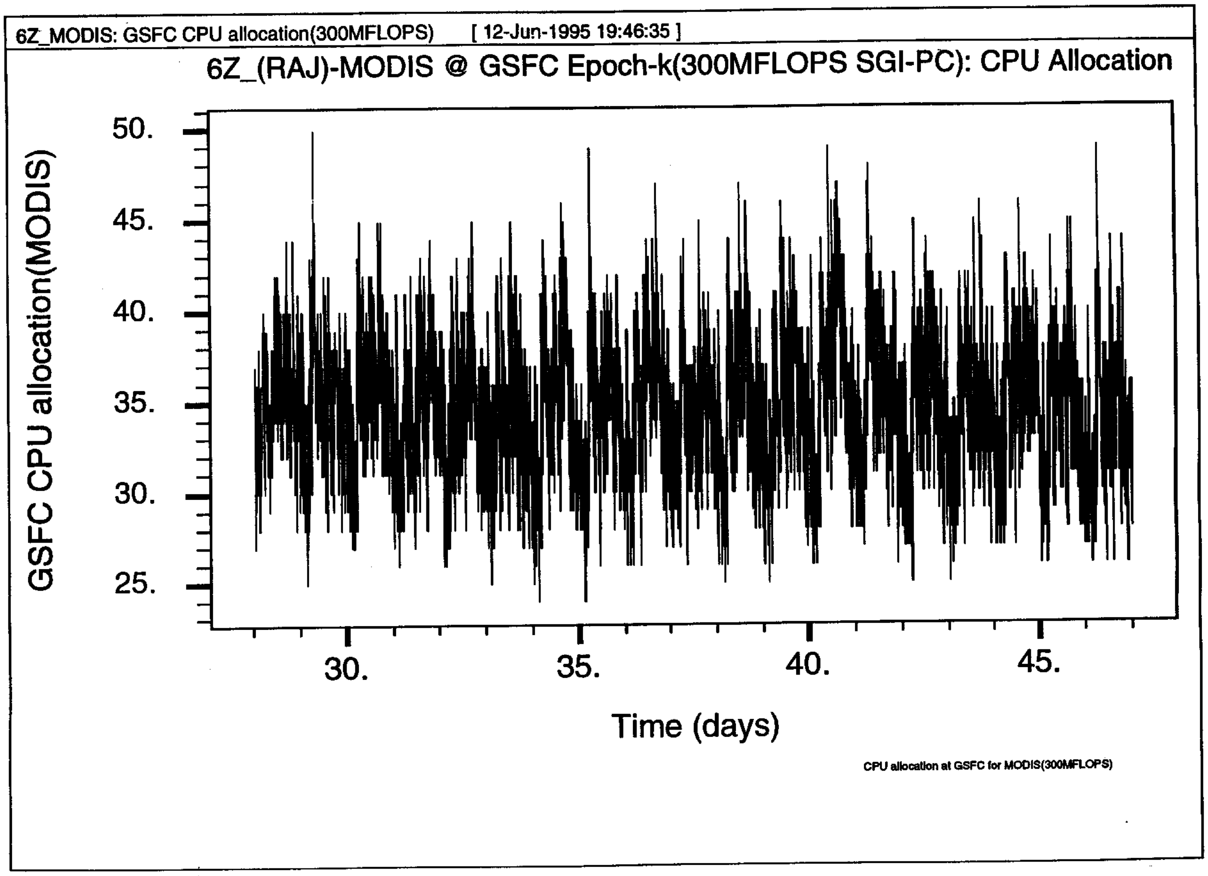

Figure 6.3-7a. MODIS (GSFC) Processing Resource Usage - Number of CPUs (300 MFLOPS)

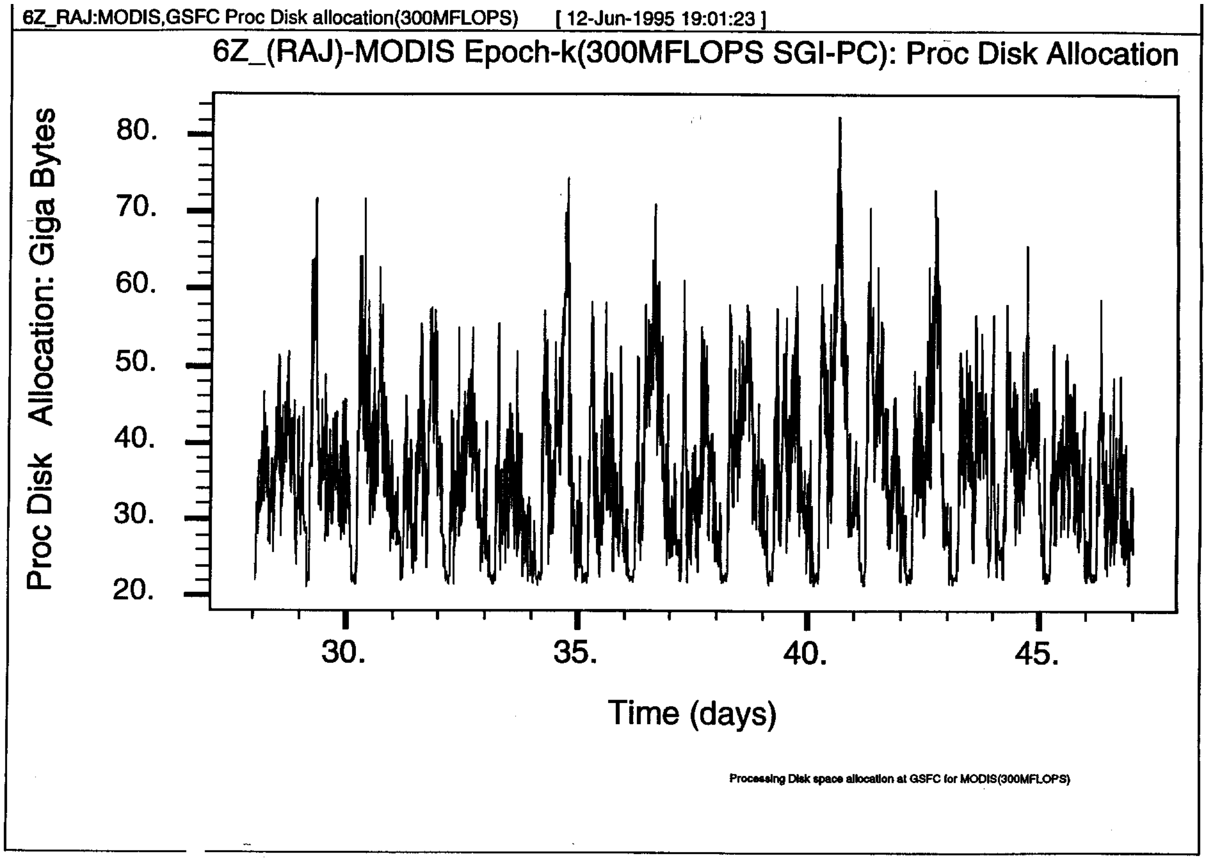

Figure 6.3-7b. MODIS (GSFC) Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

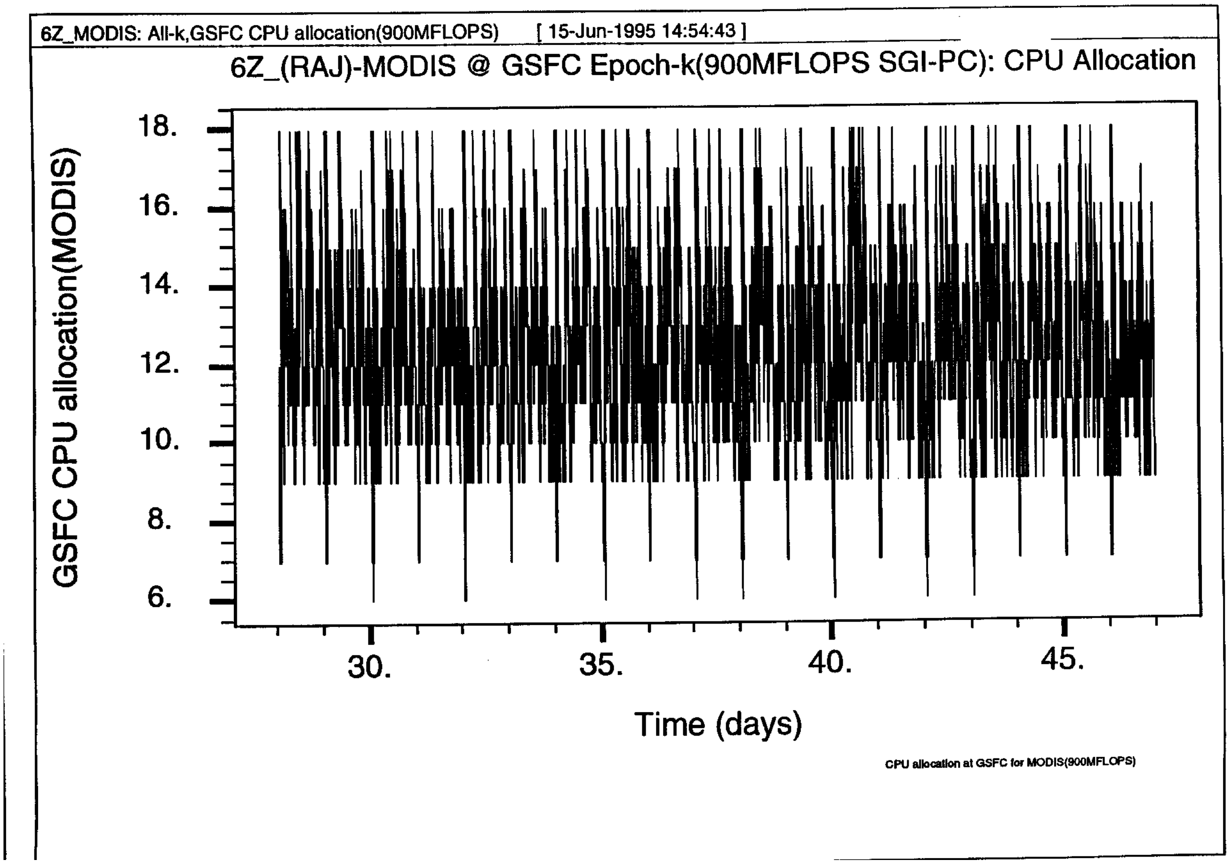

Figure 6.3-8a. MODIS (GSFC) Processing Resource Usage - Number of CPUs (900 MFLOPS)

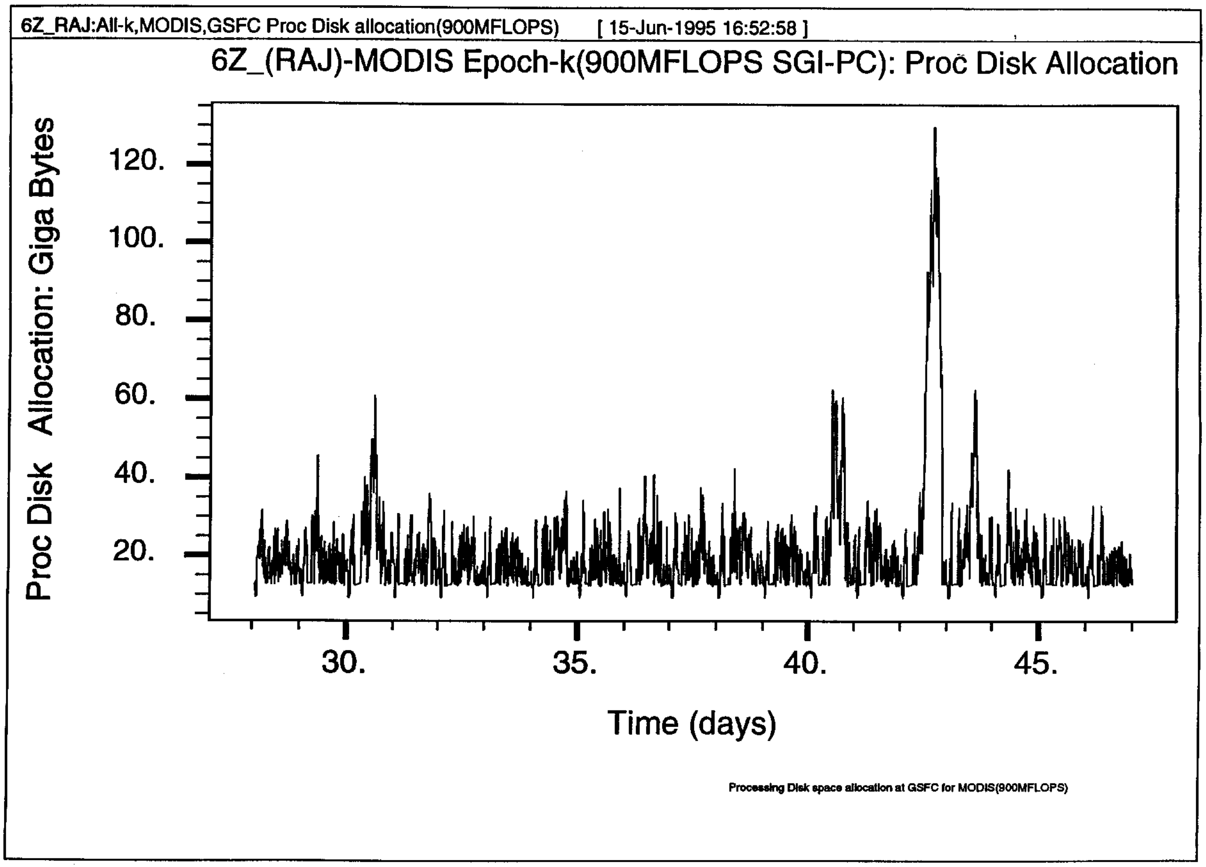

Figure 6.3-8b. MODIS (GSFC) Processing Resource Usage - Processing Disk Capacity (900 MFLOPS)

Table 6.3-24a. MODIS (GSFC) CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
34.5	53 (constrained)	35.7	83.3

Table 6.3-24b. MODIS (GSFC) CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
12.3	18 (constrained)	20.5	131.4

6.3.1.2.6 MODIS at EDC

The January 1995 version of the baseline did not contain all the information about MODIS L3 processing at EDC required by the ECS Systems Performance Model. Therefore, MODIS processing at EDC was not simulated.

6.3.1.2.7 ASTER

The ASTER processing was dynamically simulated for a period of three weeks after steady state has been attained. A 300-MFLOPS (peak) mid-range and 900-MFLOPS high-range processors with a processor efficiency factor of 0.25 were assumed for the simulation. Recall from Table 5.3-7 that ASTER processes 9 (polar cloud map) and 10 (DEM) are CPU intensive while 4 (atmospheric correction - VNIR and SWIR) and 5 (atmospheric correction -TIR) are I/O intensive. Process 4 is activated 70 times a day.

6.3.1.2.7.1 ASTER Process Completion Times

The minimum and the maximum times include job wait time before execution can occur. Tables 6.3-25a,b list the process completion times for each ASTER process run on 300- and 900-MFLOPS processors, respectively. From Table 6.3-25a, both CPU intensive jobs (Processes 9 and 10) and I/O intensive jobs (Processes 4 and 5) run for at least an hour and half on the average. The large standard deviations reflect job queue times to vary depending upon availability of all input files. Table 6.3-25b outlines process completion times when 900 MFLOPS processors are used.

Table 6.3-25a. ASTER Process Completion Times (300 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
AST_PGE_01	170.18	1.10	62.80	44.49
AST_PGE_02	187.03	4.07	104.85	48.97
AST_PGE_03	169.55	0.28	63.79	46.17
AST_PGE_04	190.67	10.27	94.13	44.24
AST_PGE_05	182.31	4.56	111.91	41.35
AST_PGE_06	183.38	0.73	73.24	43.66
AST_PGE_07	171.37	4.10	67.91	41.92
AST_PGE_08	177.63	4.04	70.36	43.61
AST_PGE_09	166.49	25.41	87.38	42.87
AST_PGE_10	235.98	117.42	178.70	41.51

Table 6.3-25b. ASTER process completion times (900 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
AST_PGE_01	142.01	0.38	42.02	28.66
AST_PGE_02	238.62	1.35	105.21	69.47
AST_PGE_03	142.29	0.28	42.85	28.98
AST_PGE_04	232.97	20.15	129.08	39.76
AST_PGE_05	238.69	37.19	156.94	42.19
AST_PGE_06	213.06	3.20	86.63	54.56
AST_PGE_07	235.62	2.05	85.70	55.25
AST_PGE_08	234.85	1.83	87.93	56.88
AST_PGE_09	138.68	27.57	70.91	32.44
AST_PGE_10	114.21	39.77	77.82	23.42

6.3.1.2.7.2 ASTER Processing Resource Usage

The dynamically simulated processing resource usage is illustrated in Figures 6.3-9a,b (number of CPUs and processing disk capacity) for 300 MFLOPS processors. Also indicated in Table 6.3-26a, are the corresponding average and peak number of CPUs and staging disk capacities. Figure 6.3-9a and Table 6.3-26a indicate that an average of 3.2 processors (with 300 MFLOPS peak) would suffice for ASTER with an average staging disk volume of approximately 27.8 GB. ASTER processing resource usage for 900 MFLOPS processors is shown in Figure 6.3-10a,b and Table 6.3-26b

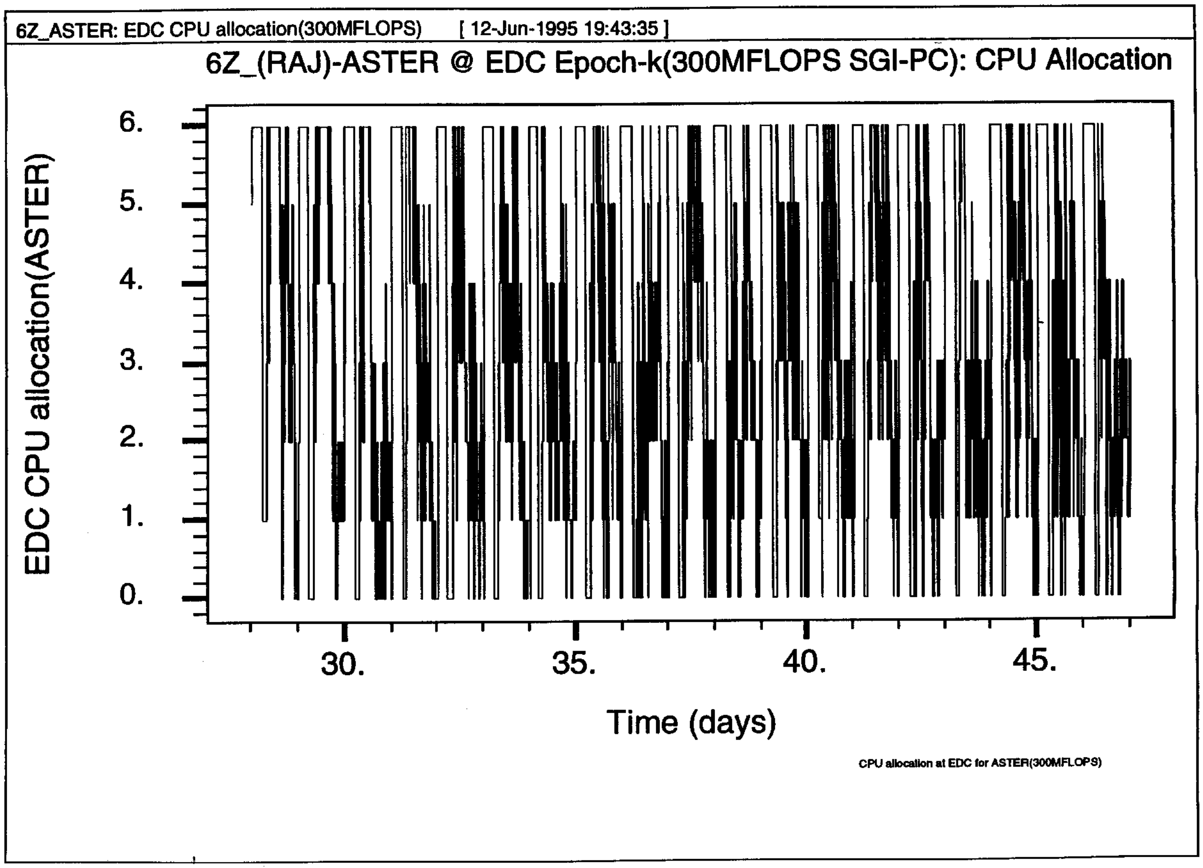

Figure 6.3-9a. ASTER Processing Resource Usage - Number of CPUs (300 MFLOPS)

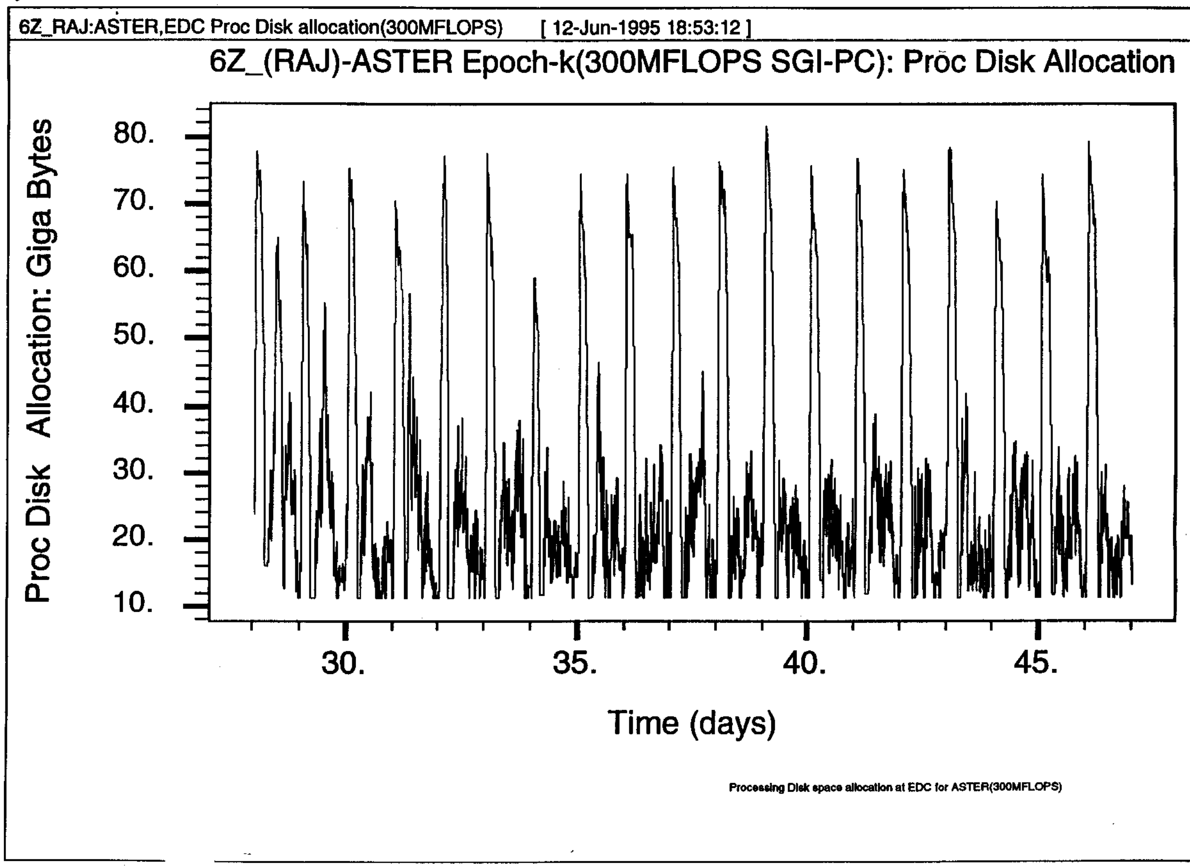

Figure 6.3-9b. ASTER Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

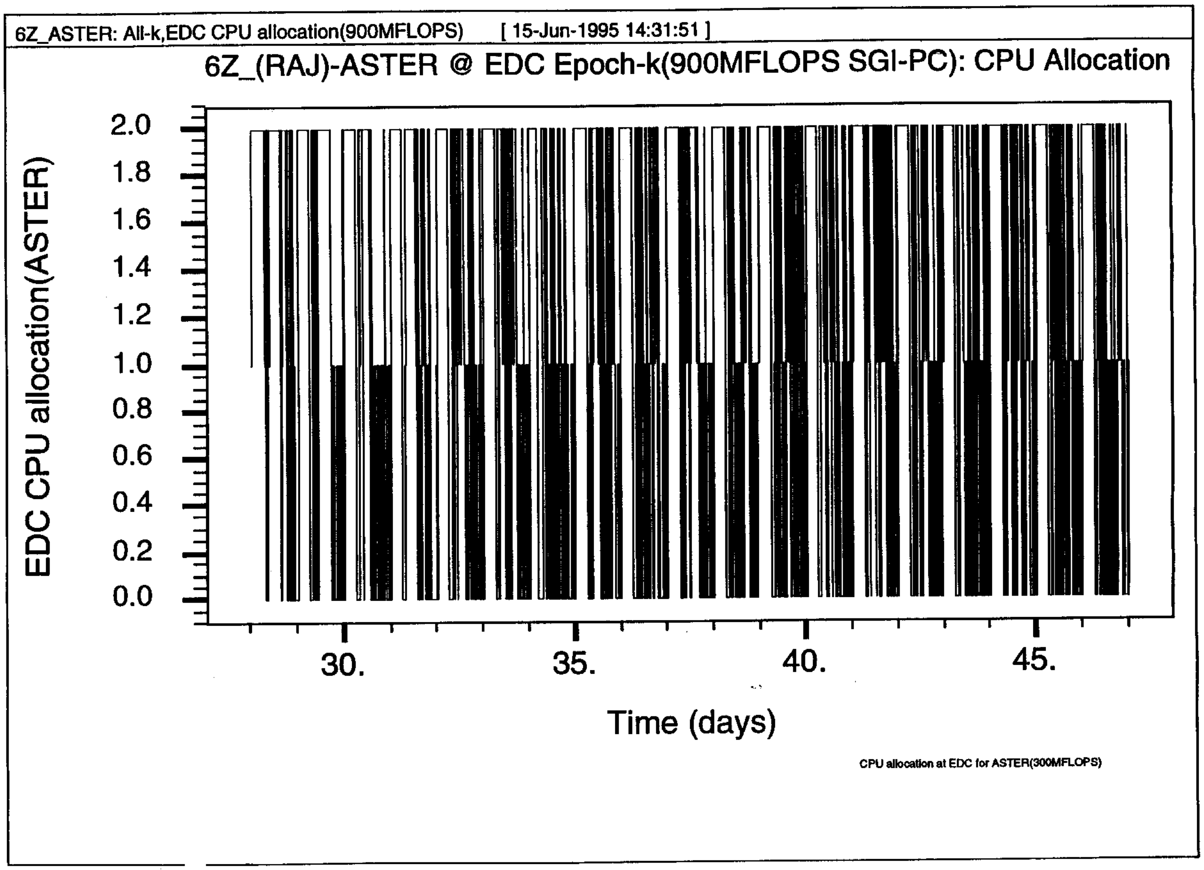

Figure 6.3-10a. ASTER Processing Resource Usage - Number of CPUs (900 MFLOPS)

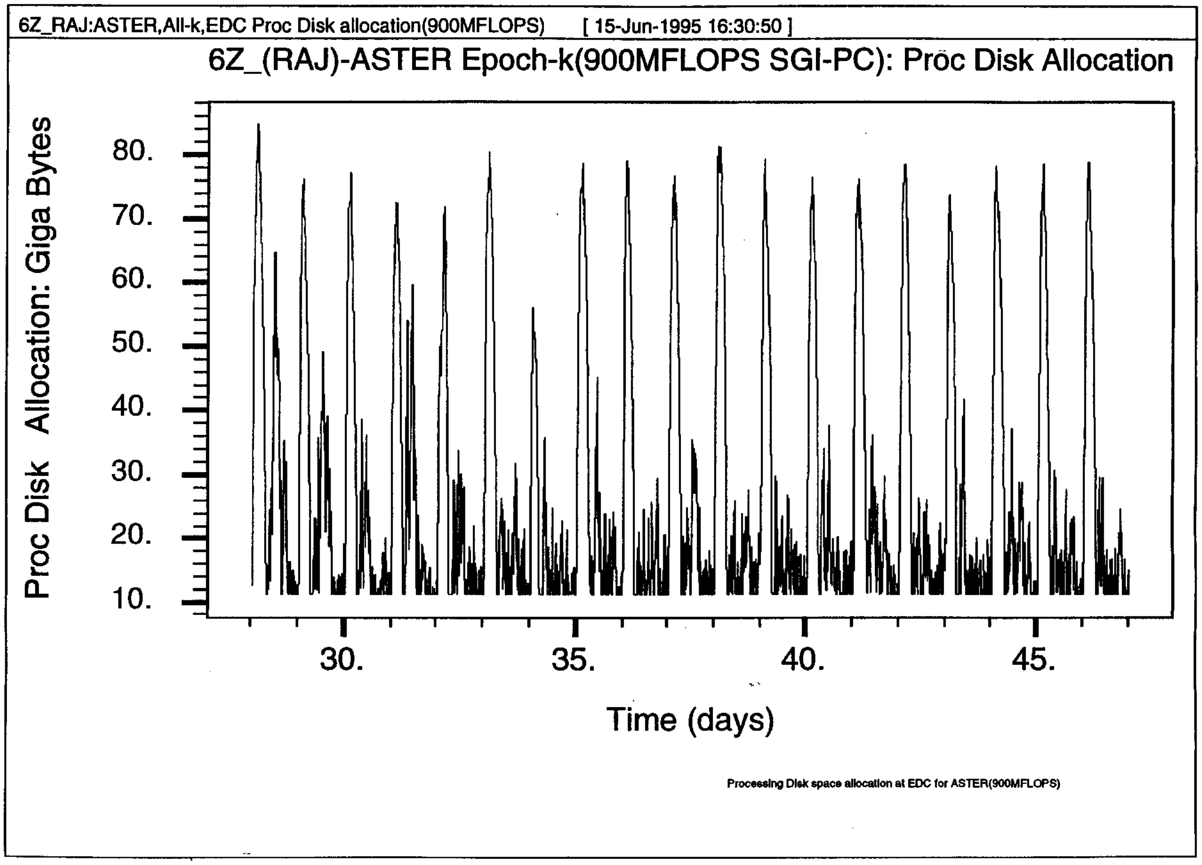

Figure 6.3-10b. ASTER Processing Resource Usage - Processing Disk Capacity (900 MFLOPS)

Table 6.3-26a. ASTER CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
3.2	6.0 (constrained)	27.8	82.1

Table 6.3-26b. ASTER CPU and Staging Disk Capacity from ECS Systems Performance Model (900 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
1.1	2.0 (constrained)	25.7	81.9

6.3.1.2.8 LIS

The LIS processing was dynamically simulated for a period of three weeks after steady state has been attained. A 300-MFLOPS (peak) mid-range processors with an efficiency factor of 0.25 were assumed for the simulation. Refer to Table 5.3-9 for LIS processing requirements.

6.3.1.2.8.1 LIS Process Completion Times

Table 6.3-27 shows the process completion times for each LIS process run on 300-MFLOPS peak processors. The minimum and the maximum times include job wait time before execution can occur. The number of processors were constrained to 4. As indicated by the process completion times, LIS processing is not intensive.

Table 6.3-27. LIS Process Completion Times (300 MFLOPS)

Processes	T_{max} (minutes)	T_{min} (minutes)	T_{avg} (minutes)	T_{stddev} (minutes)
LIS	44.96	0.55	1.98	4.78
LIS_M	0.02	0.02	0.02	0.00

6.3.1.2.8.2 LIS Processing Resource Usage

The dynamically simulated processing resource usage is illustrated in Figures 6.3-11a,b as number of processing CPUs, processing disk capacities, respectively, for 300 MFLOPS processors. Also indicated in Table 6.3-28, is the corresponding average and peak number of CPUs and staging disk capacity. Figures 6.3-11a,b and Table 6.3-28 indicate that an average of 2.1 processors (with 300 MFLOPS peak) would suffice for LIS with an average staging disk volume of approximately 1.7 GB.

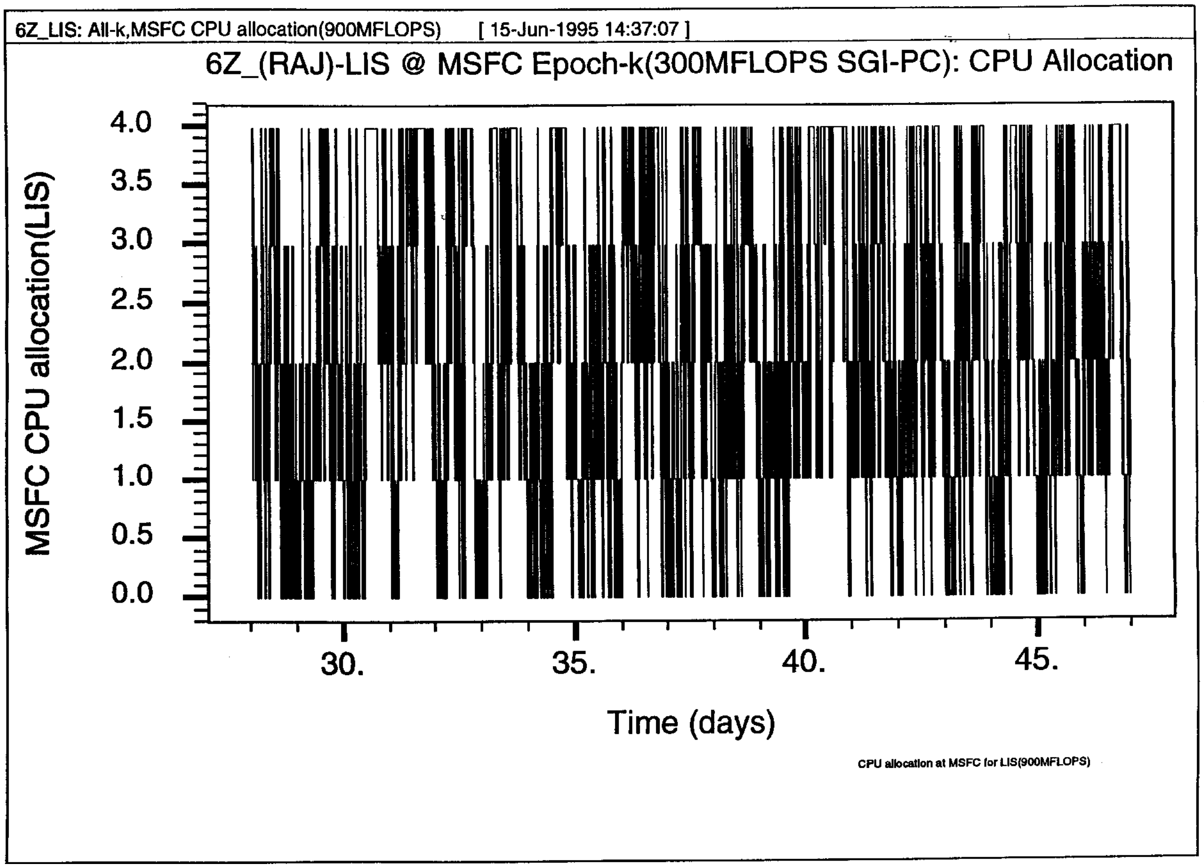

Figure 6.3-11a. LIS Processing Resource Usage - Number of CPUs (300 MFLOPS)

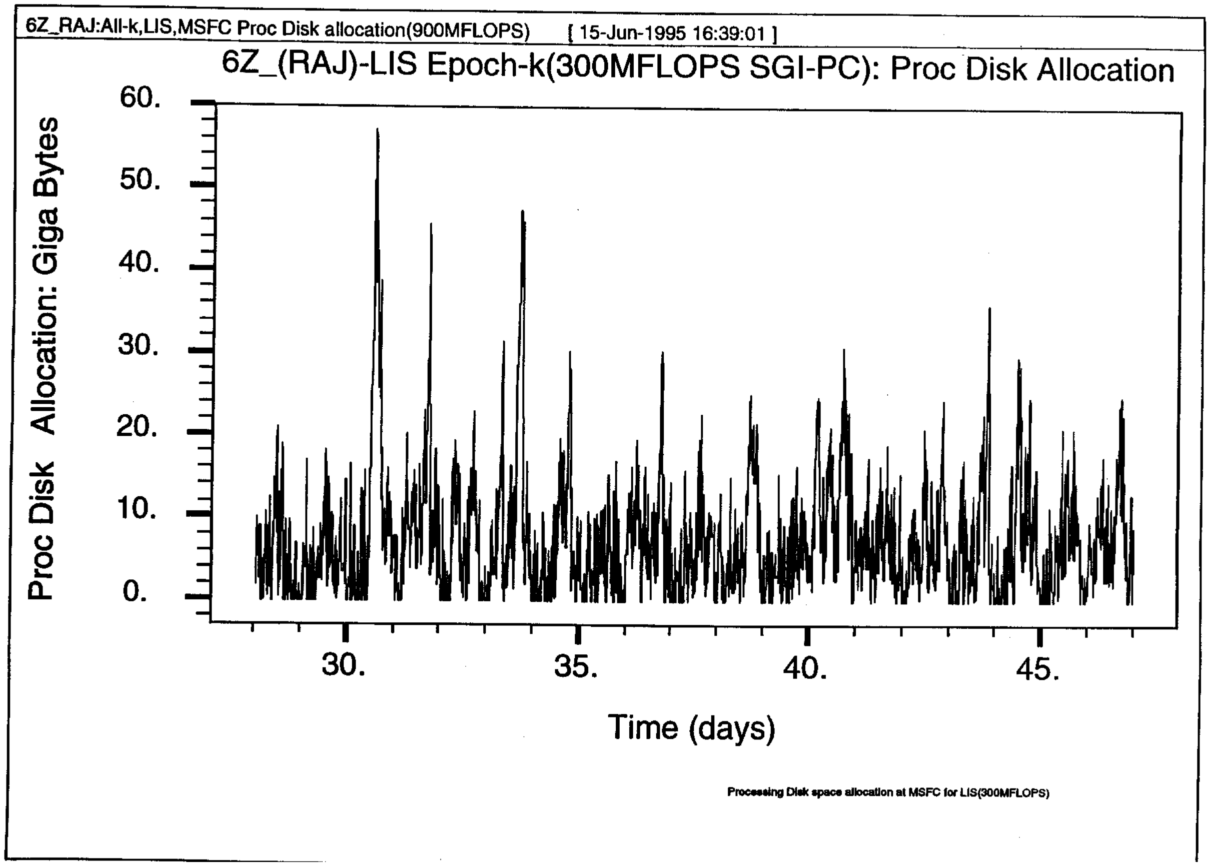

Figure 6.3-11b. LIS Processing Resource Usage - Processing Disk Capacity (300 MFLOPS)

Table 6.3-28. LIS CPU and Staging Disk Capacity from ECS Systems Performance Model (300 MFLOPS)

Number of CPUs		**Staging disk capacity (GB)**	
Average	Peak	Average	Peak
2.1	4.0 (constrained)	1.7	7.5

6.3.1.3 Processing <--> Data Handler Network Flow by DAAC

Table 6.3-29 shows the flow into the network from the processors. It is clear that GSFC and LaRC have the largest flows into the network.

Table 6.3-29. Processing <--> Data Handler Network Flow by DAAC

	Network Flow (MB/s)	
DAAC	**300 MFLOPS**	**900 MFLOPS**
EDC	4.72	4.76
GSFC	13.63	13.64
LaRC	13.67	14.19
MSFC	3.41	3.41

6.3.1.4 Advantages

- Processing each instrument on a cluster will make administration easier
- Instrument requirements can be directly mapped to hardware;
- Product chains can be easily handled. If output from a Level 1 process is input to a Level 2 process, the files can remain in the staging area until the Level 2 process is ready to run. There can be substantial savings in the cost of moving data;
- Since the algorithms come from diverse instruments, it is believed that there may be special software and hardware requirements for each instrument. In this configuration, special hardware or software requirements can be localized on one cluster.

6.3.1.5 Disadvantages

- An instrument's cluster can remain idle while other instruments' processing can have a backlog;
- Additional cluster infrastructure is necessary even if an instrument's processing requirements are small (e.g. MOPITT);
- Recovery time may be more in case of failure.

6.3.2 One instrument's products per cluster except for selected products requiring major processing resources

The requirements of each instrument (see Section 5.3) and resulting analyses (see Section 6) have yielded the following:

- It may be appropriate for MISR to be assigned to a high performance cluster at the LaRC DAAC. The number of activations per day combined with data volumes staged and destaged, and product chain dependencies may make it prohibitive to assign MISR processes to different clusters. A dynamic simulation is necessary for further understanding;
- CERES Subsystems 4 and 5 may be suitable for processing on a shared high performance cluster with MISR. Again, a dynamic simulation is necessary.

A dynamic analysis with the ECS Systems Performance Model must be performed to fully understand the suitability of this alternative for data production.

6.3.2.1 Advantages

- Special hardware/software requirements of multiple instruments can be localized.

6.3.2.2 Disadvantages

- For processing requiring major processing resources and having product chain dependencies, data may have to be moved to another cluster. This may unnecessarily increase data hops.

6.3.3 Multiple instruments' products per cluster

This may apply to conditions whereby instruments with interdependent processing may be collocated. Both CERES and MISR require MODIS products generated at GSFC. However, the MODIS products required by CERES and MISR are different. Also, there is no interdependency among the three instruments. Therefore, the LaRC scenario is not ideal for analyzing this optimization alternative. Currently there are no instrument dependencies within the same DAAC.

6.3.3.1 Advantages

- Instruments dependent on one another can be collocated, thereby, minimizing data hops.

6.3.3.2 Disadvantages

- For processing requiring major processing resources and having product chain dependencies, data may have to be moved from one cluster to another. This may unnecessarily increase data hops;
- It may be more difficult to map individual processing requirements to hardware selection;
- Each instrument's growth in processing requirements can have global consequences.
- Instrument unique hardware/software requirements cannot be localized. Duplication may be necessary that can drive up costs;
- For processes with product chain dependencies, data may have to be moved from one location to another. This may unnecessarily increase data hops;
- This alternative may optimize resources but can introduce additional complexities for the Planning and Data Processing Subsystems.

6.3.4 Any instrument's products on any cluster that can support it; selected by current processing load

This option is a mix-and-match situation. The processing load will determine the cluster where a particular instrument's data will be processed. This alternative allows the use of a large supercomputer to process many instruments at a DAAC site. A dynamic analysis with the ECS

Systems Performance Model must be performed to fully understand the suitability of this alternative for data production.

6.3.4.1 Advantages

- Idle time can be minimized because jobs are processed depending upon the current processing load on a cluster.

6.3.4.2 Disadvantages

- Mapping requirements to hardware is more difficult;
- Each instrument's growth in processing requirements can have global consequences;
- Instrument unique hardware/software requirements cannot be localized. Duplication may be necessary that can drive up costs;
- For processes with product chain dependencies, data may have to be moved from one location to another. This may unnecessarily increase data hops;
- This alternative may optimize resources but can introduce additional complexities for the Planning and Data Processing Subsystems.

7. Summary and Recommendations

Physical cluster optimization, on a site-by-site basis for release A, is not a major concern due to the small numbers and scale of the physical equipment currently envisioned for activation at that time. Release A implementations predicted for operations for LaRC and MSFC involve mid-performance (predicted) LANs and only two physical science processors within the SPRHW CI [1]. The GSFC configuration, which does not support processing operations, involves one or a small number of compute resources at a maximum. Thus, single physical subnetworks can be used (with the proper backup for RMA concerns) to couple the processing resources with primary ingest and Data Server resources, for example. The driver on selecting more than one subnetwork will be the actual throughput rates required, as opposed to operational or mission requirements which form the real basis of the implementation alternatives summarized earlier. It is expected that physical subnetwork optimization will be a larger issue for releases B and beyond. This paper is intended to provide a framework to the kind of analysis that is necessary to study various cluster optimization alternatives identified in the previous version of this paper. The first processing alternative, namely "one instrument per cluster" is analyzed in detail using the ECS Systems Performance Model. Various factors that determine the most efficient configuration are identified. It should be noted that because the ECS Systems Performance Model is evolving and has not matured to a level required for studying the other cluster optimization alternatives listed in this paper, conclusions cannot be made at this stage on the most suitable configuration for each DAAC site. However, the key recommendation is that multiple strings/cluster/subnetwork formation alternatives and selection criteria be allowed both between DAAC sites and within for Releases B and beyond. One implementation alternative for all sites and all releases is not recommended. This will permit subnetworks of ECS resources to be tuned to meet the primary needs of the DAAC site, but will not disallow the view of the resources (through planning and production management) as a single processing pool or series of subpools.

For Releases B and beyond, it is important that cluster optimization alternatives identified in this study be considered before selection of Data Processing hardware classes. The alternatives can potentially optimize communications, staging storage and ease operations management and control. A more detailed study with the ECS Performance Model is necessary and recommended to explore all cluster optimization alternatives.

Abbreviations and Acronyms

ACRIM	Active Cavity Radiometer Irradiance Monitor
AHWGP	Ad Hoc Working Group on Production
AI&T	Algorithm Integration & Test
AITHW	Algorithm Integration and Test Hardware
AQAHW	Algorithm Quality Assurance Hardware
ASTER	Advanced Spaceborne Thermal Emission and Reflection Radiometer
BONeS	Block Oriented Network Simulation
BW	Bandwidth
CDR	Critical Design Review
CERES	Clouds and Earth's Radiant Energy System
CPU	Central Processing Unit
DAAC	Distributed Active Archive Center
DEM	Digital Elevation Model
DPS	Data Processing Subsystem
ECS	EOSDIS Core System
EDC	EROS Data Center
EOSDIS	Earth Observing System Data Information System
EROS	Earth Resources Observation System
GAC	Global Area Coverage
GB	Gigabytes
GSFC	Goddard Space Flight Center
HWCI	Hardware Configuration Item
I&T	Integration & Test
I/O	Input/Output
L1	Level 1
L2	Level 2

L3	Level 3
LAN	Local Area Network
LaRC	Langley Research Center
LIS	Lightning Imaging Sensor
MB	Megabytes
MFLOPS	Millions of Floating Point Operations Per Second
MFPOs	Millions of Floating Point Operations
MISR	Multi-Angle Imaging Spectroradiometer
MODIS	Moderate-Resolution Imaging Spectroradiometer
MSFC	Marshall Space Flight Center
NAS	Network Attached Storage
NOAA	National Oceanic and Atmospheric Administration
PDR	Preliminary Design Review
PGE	Product Generation Executive
QA	Quality Assurance
RMA	Reliability, Maintainability, Availability
SDPS	Science Data Processing Segment
SDR	System Design Review
SDS	System Design Specification
SPRHW	Science Processing Hardware
TRMM	Tropical Rainfall Measuring Mission
V0	Version Zero